arXiv:1603.02982v3 [stat.ME] 19 Oct 2016

Functional Autoregression for Sparsely Sampled Data

Daniel R. Kowal, David S. Matteson, and David Ruppert*

October 20, 2016

Abstract

We develop a hierarchical Gaussian process model for forecasting and inference of functional time series data. Unlike existing methods, our approach is especially suited for sparsely or irregularly sampled curves and for curves sampled with non-negligible measurement error. The latent process is dynamically modeled as a functional autoregression (FAR) with Gaussian process innovations. We propose a fully nonparametric dynamic functional factor model for the dynamic innovation process, with broader applicability and improved computational efficiency over standard Gaussian process models. We prove finite-sample forecasting and interpolation optimality properties of the proposed model, which remain valid with the Gaussian assumption relaxed. An efficient Gibbs sampling algorithm is developed for estimation, inference, and forecasting, with extensions for FAR(p) models with model averaging over the lag p. Extensive simulations demonstrate substantial improvements in forecasting performance and recovery of the autoregressive surface over competing methods, especially under sparse designs. We apply the proposed methods to forecast nominal and real yield curves using daily U.S. data. Real yields are observed more sparsely than nominal yields, yet the proposed methods are highly competitive in both settings.

KEY WORDS: functional factor analysis; Gaussian process; hierarchical Bayes; model averaging; time series.

*Kowal is PhD Candidate, Department of Statistical Science, Cornell University, 301 Malott Hall, Ithaca, NY 14853 (E-mail: drk92@cornell.edu). Matteson is Assistant Professor, Department of Statistical Science and ILR School, Cornell University, 1196 Comstock Hall, Ithaca, NY 14853 (E-mail: matteson@cornell.edu; Webpage: `http://www.stat.cornell.edu/~matteson/`). Ruppert is Andrew Schultz, Jr. Professor of Engineering, Department of Statistical Science and School of Operations Research and Information Engineering, Cornell University, 1170 Comstock Hall, Ithaca, NY 14853 (E-mail: dr24@cornell.edu; Webpage: `http://people.orie.cornell.edu/~davidr/`). Financial support from NSF grant AST-1312903 (Kowal and Ruppert) and a Xerox PARC Faculty Research Award, and NSF grant DMS-1455172 (Matteson) is gratefully acknowledged. We thank Professor Robert A. Jarrow for suggesting the yield curve application and for his helpful discussions. The authors thank the editors and two referees for very helpful comments.

1 Introduction

We develop a hierarchical Gaussian process model for forecasting and inference of functional time series data. A *functional time series* is a time-ordered sequence of random functions, $Y_1, \ldots, Y_T$, on some compact index set $\mathcal{T} \subset \mathbb{R}^D$, typically with $D = 1$. Unlike existing methods, our approach is especially suited for sparsely or irregularly sampled curves, in which the functions $Y_t(\tau)$ are observed at a small number of possibly unequally-spaced points $\tau \in \mathcal{T}$, and for curves sampled with non-negligible measurement error, which occur frequently in financial applications. Applications of functional time series are abundant, including: daily or weekly interest rate curves as a function of time to maturity, such as daily Eurodollar futures contracts (Kargin and Onatski, 2008) and weekly yield curves (Hays et al., 2012; Kowal et al., 2016); yearly sea surface temperature as a function of time-of-year (Besse et al., 2000); yearly mortality and fertility rates as a function of age (Hyndman and Ullah, 2007); daily pollution curves as a function of time-of-day (Damon and Guillas, 2002; Aue et al., 2015); and a vast collection of spatio-temporal applications in which a time-dependent variable is measured as a function of spatial location (e.g., Cressie and Wikle, 2011). The primary goal of functional time series analysis is usually forecasting $\{Y_t\}$, but we are also interested in performing inference and obtaining an interpretable representation of the time evolution of $\{Y_t\}$.

The most prevalent model for functional time series data is the *functional autoregressive model of order 1*, written FAR(1):

$$Y_t - \mu = \Psi(Y_{t-1} - \mu) + \epsilon_t, \tag{1}$$

where $Y_t \in L^2(\mathcal{T})$, Ψ is a bounded linear operator on $L^2(\mathcal{T})$, $\epsilon_t \in L^2(\mathcal{T})$ is a sequence of independent mean zero random innovation functions with $\mathbb{E}||\epsilon_t||^2 < \infty$, and μ is the mean of $\{Y_t\}$ under stationarity. The FAR(1) model, developed by Bosq (2000), is an extension of two highly successful models: the functional linear model for function-on-function regression

and the vector autoregressive model for multivariate time series, and has been successfully applied in a variety of applications. Importantly, the FAR(1) model provides a mechanism for modeling the evolution of $\{Y_t\}$ jointly over the entirety of the domain $\mathcal{T}$. More generally, (1) can be extended for multiple lags to the FAR(p) model: $Y_t - \mu = \sum_{\ell=1}^{p} \Psi_\ell(Y_{t-\ell} - \mu) + \epsilon_t$.

Existing approaches for estimating the FAR(p) model typically use an eigendecomposition of the empirical (contemporaneous and lagged) covariance operators (Damon and Guillas, 2002, 2005; Horváth and Kokoszka, 2012; Kokoszka, 2012) or kernel-based procedures for modeling the conditional expectation (Besse et al., 2000). A related approach is to estimate a multivariate time series model for the functional principal component (FPC) scores of the observed data (Aue et al., 2015). Extensions of the FAR(1) model for nonstationary functional time series are available, such as the time-dependent FAR kernels proposed in Chen and Li (2015).

In general, existing methods for FAR(p) are designed for functional data observed on dense grids without measurement error, and typically require pre-smoothing discretized functional observations. However, such procedures may exhibit erratic behavior for sparse designs and are inappropriate in such settings. More generally, under an FAR(p) model that includes measurement error and discretization of the functional observations, we prove that the two most common approaches for functional data analysis—estimators that are linear in the FPC scores or the pre-smoothed observations—produce predictions that are inadmissible (in a decision theory sense). Indeed, the presence of measurement error fundamentally alters the behavior of the observable process: if an FAR process is observed with measurement error, then the observable process is no longer an FAR process, but rather a functional autoregressive moving average process (see Proposition 1). Even under dense designs, existing methods produce poor estimates of the FAR operator Ψ (Didericksen et al., 2012), which inhibits interpretability of the time evolution of $\{Y_t\}$, and do not provide finite-sample inference. We propose new methodology that simultaneously addresses all of these challenges.

We propose a general two-level hierarchy for modeling functional time series: an *observa-*

tion equation addresses measurement error and discretization of the functional data, while an *evolution equation* defines a process model for the underlying functional time series. The latent process is dynamically modeled as an FAR(p). We parsimoniously specify the FAR model with mean zero Gaussian process innovations, which are fully specified by covariance functions without parameterizing sample paths. The dynamic innovation process is further specified by a dynamic functional factor model. In contrast with standard approaches for Gaussian processes, this avoids selecting and estimating a parametric covariance function, and allows greater computational stability and efficiency, and broader applicability. Interpolating curves at unsampled locations and forecasting future curves are primary objectives in functional time series modeling; the proposed model produces optimal (best linear) predictions under both sparse and dense designs in the presence of measurement error, even with the Gaussian assumption relaxed. We propose an efficient Gibbs sampling algorithm for estimation, inference, and forecasting. Extensive simulations demonstrate substantial improvements in forecasting performance and recovery of the autoregressive surface over competing methods, especially under sparse designs.

We apply our methodology to model and forecast *nominal* and *real* yield curves using daily U.S. data. For a given currency and level of risk of a debt, the nominal yield curve, $Y_t^N(\tau)$, describes the interest rate at time t as a function of the length of the borrowing period, or time to maturity, τ. Similarly, the real yield curve, $Y_t^R(\tau)$, corresponds to an interest rate that is adjusted for inflation. Both Y_t^N and Y_t^R may be modeled as functional time series. However, real yields are sparsely observed for each time t, and only at longer maturities, which is problematic for existing functional time series models. The proposed methods provide a natural hierarchical framework for modeling both nominal yield curves and real yield curves, and in both cases produce highly competitive forecasts.

Bayesian methods for functional time series are limited, with the exception of Laurini (2014) and Kowal et al. (2016). The primary contributions of this article are the following: (i) development of a hierarchical framework for FAR(p) (Section 2), which produces optimal

(best linear) predictions under both sparse and dense designs in the presence of measurement error; (ii) a dynamic functional factor model for the innovation covariance, which is non-parametric, computationally convenient, and offers useful generalizations to non-Gaussian distributions (Section 3); (iii) a procedure for model averaging over the lag, p, within a hierarchical FAR(p) model (Section 4); (iv) comparisons of the proposed methods to existing methods for FAR(p) using theoretical results (Section 5), an extensive simulation study (Section 6), and a real data application (Section 7); (v) a comparative forecasting study of daily U.S. nominal and real yield curve data (Section 7); and (vi) an efficient Gibbs sampling algorithm, which uses common full conditional distributions and existing `R` software (Appendix). Details of our Gibbs sampling algorithm and additional theoretical and simulation results are in the web supplement.

2 Hierarchical Gaussian Processes for FAR

Let $Y_1, \ldots, Y_T$ be a time-ordered sequence of random functions in $L^2(\mathcal{T})$, where $\mathcal{T} \subset \mathbb{R}^D$ is a compact index set. We focus on $D = 1$ with $\mathcal{T} = [0, 1]$, but the methods can be developed more generally. For interpretability and computational convenience, we restrict our attention to the integral operators defined by $\Psi_\ell(Y)(\tau) = \int \psi_\ell(\tau, u) Y(u)\, du$, so the FAR($p$) model is

$$Y_t(\tau) - \mu(\tau) = \sum_{\ell=1}^{p} \int \psi_\ell(\tau, u) \left\{ Y_{t-\ell}(u) - \mu(u) \right\} du + \epsilon_t(\tau) \quad \forall \tau \in \mathcal{T}. \tag{2}$$

Using integral operators, the FAR(p) model resembles the functional linear model, in which $(Y_t - \mu)$ is regressed on $(Y_{t-1} - \mu), \ldots, (Y_{t-p} - \mu)$. The functional linear model is widely popular in functional data analysis, and has been extensively studied (e.g., Cardot et al., 1999; Ramsay, 2006).

In practice, model (2) is incomplete: the functional observations $\{Y_t\}$ are not observed directly, but rather via discrete samples of each curve, and typically with measurement error.

Suppose that we observe $y_{i,t} \in \mathbb{R}$ sampled with noise $\nu_{i,t}$ from $Y_t \in L^2(\mathcal{T})$:

$$y_{i,t} = Y_t(\tau_{i,t}) + \nu_{i,t} \tag{3}$$

for $i = 1, \ldots, m_t$, where $\tau_{1,t}, \ldots, \tau_{m_t,t}$ are the observation points of Y_t and $\nu_{i,t}$ is a mean zero measurement error with finite variance. Typically for functional data, m_t will be large and $\mathcal{T}_t = \{\tau_{1,t}, \ldots, \tau_{m_t,t}\}$ will be dense in $\mathcal{T}$. However, for our procedures, we allow m_t to be small for some (or all) t, with observation points $\mathcal{T}_o \equiv \cup_t \mathcal{T}_t$ dense or sparse in $\mathcal{T}$. Combining (3) with (2) for $p = 1$ and defining $\mu_t \equiv Y_t - \mu$, we obtain the two-level hierarchical model

$$\begin{cases} y_{i,t} = \mu(\tau_{i,t}) + \mu_t(\tau_{i,t}) + \nu_{i,t}, & i = 1, \ldots, m_t, \\ \mu_t(\tau) = \int \psi(\tau, u)\mu_{t-1}(u)\, du + \epsilon_t(\tau), & \forall \tau \in \mathcal{T} \end{cases} \tag{4}$$

for $t = 2, \ldots, T$, where we assume that $\{\nu_{i,t}\}$ and $\{\epsilon_t\}$ are mutually independent sequences.

The measurement error is a nontrivial component of model (4), which we demonstrate in the following proposition:

Proposition 1. *Let* $Y_t - \mu = \sum_{\ell=1}^{p} \Psi_\ell(Y_{t-\ell} - \mu) + \epsilon_t$*, and suppose that we observe* $y_t = Y_t + \nu_t$*, where* $\{\epsilon_t\}$ *and* $\{\nu_t\}$ *are independent white noise processes. Then the observable process* $\{y_t\}$ *follows a* functional autoregressive moving average *(FARMA) process of order* (p, p)*.*

We define a FARMA process and prove Proposition 1 in Section B.1 of the web supplement. The implication of Proposition 1 is that, if the true model for Y_t is FAR(p), yet Y_t is observed with error, then the FAR(p) model for the observables is inappropriate. As a result, estimation of Ψ_ℓ will be inefficient and forecasting will deteriorate, due to both increased estimation error of Ψ_ℓ and model misspecification. By comparison, the hierarchical model decomposes the observed data into a functional (autoregressive) process and measurement error, and in doing so circumvents the model misspecification issues implied by Proposition 1.

We model the random functions μ, ψ, and $\{\epsilon_t\}$ as *Gaussian processes*: $\mu \sim \mathcal{GP}(0, K_\mu)$,

$\psi \sim \mathcal{GP}(0, K_\psi)$, and $\epsilon_t \overset{indep}{\sim} \mathcal{GP}(0, K_\epsilon)$, where the notation $\mathcal{GP}(m, K)$ denotes a Gaussian process with mean function m and covariance function K. Gaussian processes have a long history in machine learning (Rasmussen and Williams, 2006) and spatial statistics (Cressie and Wikle, 2011), and have seen increased application in functional data analysis, especially for hierarchical modeling (Behseta et al., 2005; Kaufman et al., 2010; Shi and Choi, 2011; Earls and Hooker, 2014). The conditional distribution of $\mu_t = Y_t - \mu$ is $[\mu_t|\mu_{t-1}, \psi, K_\epsilon] \sim \mathcal{GP}(\int \psi(\cdot, u)\mu_{t-1}(u)\, du, K_\epsilon)$, which models the evolution of μ_t and serves as the prior distribution for the observation level of (4). Notably, the model only requires *conditionally* Gaussian processes, and therefore may accommodate more general distributional assumptions, such as scale-mixtures of Gaussian distributions and stochastic volatility. Moreover, the posterior expectations derived from the hierarchical Gaussian process model are best linear predictors, and therefore are optimal among linear predictors for interpolation and forecasting of Y_t, even for non-Gaussian distributions (see Section 5). We assume $\nu_{i,t} \overset{iid}{\sim} N(0, \sigma_\nu^2)$ for the measurement errors; priors for σ_ν^2 and the parameters associated with K_μ, K_ϵ, and K_ψ will be discussed later.

2.1 Dynamic Linear Models for FAR(p)

For practical implementation of model (4), we must select a finite set of evaluation points, $\mathcal{T}_e \equiv \{\tau_1, \ldots, \tau_M\} \subset \mathcal{T}$, at which we wish to estimate, forecast, or perform inference on the random functions, in particular $\mu_t = Y_t - \mu$. Naturally, we assume that $\mathcal{T}_t \subseteq \mathcal{T}_e$ for all t, but this assumption may be relaxed. Notably, $\mathcal{T}_e$ provides a convenient structure for forecasting and inference of $y_{i,t}$ and $Y_t(\tau_{i,t})$ at the observations points $\tau_{i,t} \in \mathcal{T}_t$, as well as interpolation of Y_t at any unobserved points, $\tau^* \in \mathcal{T}_e \setminus \mathcal{T}_o$. By definition, for any Gaussian process $x \sim \mathcal{GP}(m, K)$ defined on $\mathcal{T}$, we have $\boldsymbol{x} \sim N(\boldsymbol{m}, \boldsymbol{K})$, where $\boldsymbol{x} = (x(\tau_1), \ldots, x(\tau_M))'$, $\boldsymbol{m} = (m(\tau_1), \ldots, m(\tau_M))'$, and $\boldsymbol{K} = \{K(\tau_i, \tau_k)\}_{i,k=1}^M$. This result is particularly useful for constructing an estimation procedure and deriving the optimality results of Section 5.

By selecting M large and $\mathcal{T}_e$ dense in $\mathcal{T}$, we can accurately approximate the integral in (4) using quadrature methods:

$$\int \psi(\tau, u)\mu_{t-1}(u)\, du \approx (\psi(\tau, \tau_1), \ldots, \psi(\tau, \tau_M))\, \boldsymbol{Q}\boldsymbol{\mu}_{t-1}, \tag{5}$$

where $\boldsymbol{Q}$ is a known quadrature weight matrix and $\boldsymbol{\mu}_{t-1} = (\mu_{t-1}(\tau_1), \ldots, \mu_{t-1}(\tau_M))'$. The approximation in (5) is important for computational tractability in estimation of both μ_t and ψ. Practical implementations of functional data methods require discretization or finite approximations; the quadrature approximation in (5) is a natural approach, and does not impose restrictive assumptions on the functional forms of ψ and μ_{t-1}. In addition, our simulation analysis suggests that the quadrature approximation does not noticeably inhibit estimation or forecasting, especially relative to existing FAR methods. In practice, the trapezoidal rule for computing $\boldsymbol{Q}$ works well, and for simulated data $M = 20$ is sufficiently large. We include a sensitivity analysis in the web supplement to assess the effects of M on the approximation error in (5), which supports this choice of M.

Assuming $\mathcal{T}_o \subseteq \mathcal{T}_e$, let $\boldsymbol{Z}_t$ be the $m_t \times M$ incidence matrix that identifies the observations points observed at time t, i.e., $(\tau_{1,t}, \ldots, \tau_{m_t,t})' = \boldsymbol{Z}_t(\tau_1, \ldots, \tau_M)'$. We can write the hierarchical model (4) as a *dynamic linear model* (DLM; West and Harrison, 1997) in $\boldsymbol{\mu}_t$:

$$\begin{cases} \boldsymbol{y}_t = \boldsymbol{Z}_t\boldsymbol{\mu} + \boldsymbol{Z}_t\boldsymbol{\mu}_t + \boldsymbol{\nu}_t, & [\boldsymbol{\nu}_t|\sigma^2_\nu] \overset{indep}{\sim} N\left(\boldsymbol{0}, \sigma^2_\nu \boldsymbol{I}_{m_t}\right) \text{ for } t = 1, \ldots, T, \\ \boldsymbol{\mu}_t = \boldsymbol{\Psi}\boldsymbol{Q}\boldsymbol{\mu}_{t-1} + \boldsymbol{\epsilon}_t, & [\boldsymbol{\epsilon}_t|\boldsymbol{K}_\epsilon] \overset{indep}{\sim} N\left(\boldsymbol{0}, \boldsymbol{K}_\epsilon\right) \text{ for } t = 2, \ldots, T, \\ \boldsymbol{\mu}_1 \sim N(\boldsymbol{0}, \boldsymbol{K}_\epsilon), \end{cases} \tag{6}$$

where $\boldsymbol{y}_t = (y_{1,t}, \ldots, y_{m_t,t})'$, $\boldsymbol{\mu} = (\mu(\tau_1), \ldots, \mu(\tau_M))'$, $\boldsymbol{\Psi} = \{\psi(\tau_i, \tau_k)\}_{i,k=1}^M$, and $\boldsymbol{K}_\epsilon = \{K_\epsilon(\tau_i, \tau_k)\}_{i,k=1}^M$. Model (6) can be extended for multiple lags to the FAR(p) model by replacing the second level with $\boldsymbol{\mu}_t = \sum_{\ell=1}^p \boldsymbol{\Psi}_\ell\boldsymbol{Q}\boldsymbol{\mu}_{t-\ell} + \boldsymbol{\epsilon}_t$ for $\boldsymbol{\Psi}_\ell = \{\psi_\ell(\tau_i, \tau_k)\}_{i,k=1}^M$. The DLM formulation of the FAR(p) is useful for MCMC sampling, since efficient samplers exist for the vector-valued state variables, $\{\boldsymbol{\mu}_t\}$ (e.g., Durbin and Koopman, 2002). The pro-

posed Gibbs sampling algorithm for model (6) (see Section A of the web supplement) is a moderate extension of traditional DLM samplers, and iteratively samples the state vectors $\{\boldsymbol{\mu}_t\}$, the measurement error variance σ_ν^2, the innovation covariance $\boldsymbol{K}_\epsilon$, and the unknown evolution matrix $\boldsymbol{\Psi}$. The DLM also facilitates non-Bayesian parameter estimation and forecasting, such as an EM algorithm for the latent state variables $\{\boldsymbol{\mu}_t\}$ with the parameters $\{\sigma_\nu^2, \boldsymbol{K}_\epsilon, \boldsymbol{\Psi}\}$ (e.g., Cressie and Wikle, 2011).

The connection between the hierarchical FAR model (4) and the DLM (6) is further illuminated by considering the autocovariance properties of the respective models. Recalling $\mu_t(\tau) = Y_t(\tau) - \mu(\tau)$, let $C_\ell(\tau_1, \tau_2) = \mathbb{E}\left[\mu_t(\tau_1)\mu_{t-\ell}(\tau_2)\right]$ be the lag-ℓ autocovariance function of $\{Y_t\}$, which is time-invariant under stationarity of $\{Y_t\}$. Under model (4) and assuming stationarity of $\{Y_t\}$, the lag-1 autocovariance function is equivalently $C_1(\tau_1, \tau_2) = \mathbb{E}\left[\mu_t(\tau_1)\mu_{t-1}(\tau_2)\right] = \mathbb{E}\left[\left\{\int \psi(\tau_1, u)\mu_{t-1}(u)\, du + \epsilon_t(\tau_1)\right\}\mu_{t-1}(\tau_2)\right] = \int \psi(\tau_1, u)C_0(u, \tau_2)\, du$. For $\ell \geq 1$, we have the more general recursion $C_\ell(\tau_1, \tau_2) = \int \psi(\tau_1, u)C_{\ell-1}(u, \tau_2)\, du$, from which it is clear that each C_ℓ is completely determined by the pair (ψ, C_0). Now let $\boldsymbol{C}_\ell = \mathbb{E}\left[\boldsymbol{\mu}_t\boldsymbol{\mu}_{t-\ell}'\right]$ be the lag-ℓ autocovariance matrix for the vector-valued time series $\{\boldsymbol{\mu}_t\}$ in (6). Under stationarity of $\{\boldsymbol{\mu}_t\}$, the lag-1 autocovariance matrix of $\boldsymbol{\mu}_t$ is $\boldsymbol{C}_1 = \mathbb{E}\left[\boldsymbol{\mu}_t\boldsymbol{\mu}_{t-1}'\right] = \mathbb{E}\left[\{\boldsymbol{\Psi}\boldsymbol{Q}\boldsymbol{\mu}_{t-1} + \boldsymbol{\epsilon}_t\}\boldsymbol{\mu}_{t-1}'\right] = \boldsymbol{\Psi}\boldsymbol{Q}\boldsymbol{C}_0$. Notably, the relationship $\boldsymbol{C}_1 = \boldsymbol{\Psi}\boldsymbol{Q}\boldsymbol{C}_0$ is an approximation to the continuous version, $C_1(\tau_1, \tau_2) = \int \psi(\tau_1, u)C_0(u, \tau_2)\, du$, using the same quadrature approximation as in (5). More generally, the matrix recursion $\boldsymbol{C}_\ell = \boldsymbol{\Psi}\boldsymbol{Q}\boldsymbol{C}_{\ell-1}$ is a quadrature-based approximation to the continuous recursion, $C_\ell(\tau_1, \tau_2) = \int \psi(\tau_1, u)C_{\ell-1}(u, \tau_2)\, du$ for $\ell \geq 1$. Therefore, the evolution matrix $\boldsymbol{\Psi}\boldsymbol{Q}$ in the DLM (6) induces a discrete approximation to the autocovariance structure in the hierarchical FAR model (4).

The evolution equation of (6) resembles a VAR(1) on $\boldsymbol{\mu}_t = (\mu_t(\tau_1), \ldots, \mu_t(\tau_M))'$, but differs from a standard VAR on $\boldsymbol{y}_t$ for a few critical reasons. First, fitting a VAR to $\boldsymbol{y}_t$ is only well-defined if both the dimension m_t and the observation points $\mathcal{T}_t$ are fixed over time. If this does not hold, then imputation is necessary. Our procedure imputes automatically and optimally using the conditional mean function and the conditional covariance function

of the corresponding Gaussian process. Second, the components of $\boldsymbol{y}_t$ are likely highly correlated due to the functional nature of the observations. Strong collinearity in VARs can cause overfitting and adversely affect forecasting and inference. In our model, the kernel function ψ is regularized using a smoothness prior (see Section 4), which mitigates the adverse effects of collinearity on estimation of ψ. The smoothness prior on ψ is a nonstandard regularization technique for VARs, but is appropriate in this setting. Finally, the quadrature matrix, $\boldsymbol{Q}$, is absorbed into the VAR coefficient matrix $\boldsymbol{\Psi Q}$, and reweights the vector $\boldsymbol{\mu}_{t-1}$ using information from the evaluation points $\mathcal{T}_e$. This reweighting incorporates not only the vector values $\boldsymbol{\mu}_t$, but also the information that the components of $\boldsymbol{\mu}_t$ correspond to ordered elements of $\mathcal{T}_e$, which need not be equally spaced. The simulations of Section 6 demonstrate the substantial improvements in forecasting of our procedure relative to a VAR on $\boldsymbol{y}_t$.

3 A Dynamic Functional Factor Model for the Innovation Process

The standard approach for Gaussian process models is to select a parametric covariance function that only depends on a few parameters, and then estimate those parameters using either fully Bayesian methods or empirical Bayes (Rasmussen and Williams, 2006). The choice of the covariance function determines the properties of the sample trajectories, such as smoothness and periodicity, but notably does *not* imply a parametric form for the sample trajectories. Indeed, the FAR(1) model (6) may be estimated using these standard approaches; we provide one implementation in Section 6.

However, there are substantial computational limitations that accompany standard parametric covariance functions. Even when the covariance function is known up to some parameters $\boldsymbol{\rho}$, in general we cannot directly sample from the full conditional posterior distribution for $\boldsymbol{\rho}$. As a result, posterior sampling for $\boldsymbol{\rho}$ can be inefficient. Gaussian processes also require computation of the $M \times M$ innovation covariance matrix $\boldsymbol{K}_\epsilon$, which must be inverted—both for evaluating the conditional likelihood of $\boldsymbol{\rho}$ and for sampling $\{\boldsymbol{\mu}_t\}$ and ψ. Most common

choices for parametric covariance functions do not offer any simplifying structure for computing this inverse, which may be computationally inefficient and unstable. In addition, extensions for time-dependent covariance functions or non-Gaussian distributions are not readily available, and further increase the difficulties with posterior sampling.

We propose a low-rank, fully nonparametric approach for modeling the innovation covariance function. Using the *functional dynamic linear model* (FDLM) of Kowal et al. (2016), we estimate the unknown covariance function using a functional factor model, which does not require specification of a parametric form for the covariance function. This method avoids the need for inversion of the full $M \times M$ covariance matrix, and is more computationally stable and efficient. The integration of the FDLM into (6) retains the fully Bayesian hierarchical structure, and permits joint inference for all parameters via an efficient MCMC sampling algorithm. A *functional* factor model is most appropriate because ϵ_t is a Gaussian process with covariance function K_ϵ, so K_ϵ must be well-defined on $\mathcal{T} \times \mathcal{T}$. Notably, the FDLM offers convenient generalizations for stochastic volatility models (Kim et al., 1998) and more robust models using scale-mixtures of Gaussian distributions (Fernandez and Steel, 2000).

The FDLM decomposes the innovations ϵ_t into *factor loading curves* (FLCs), $\phi_j \in L^2(\mathcal{T})$, and time-dependent *factors*, $e_{j,t} \in \mathbb{R}$, for $j = 1, \ldots, J_\epsilon$:

$$\epsilon_t(\tau) = \sum_{j=1}^{J_\epsilon} e_{j,t}\phi_j(\tau) + \eta_t(\tau) \quad \forall \tau \in \mathcal{T}, \tag{7}$$

where J_ϵ is the number of factors and $\{\eta_t\}$ is the mean zero approximation error with $\eta_t \overset{iid}{\sim} \mathcal{GP}(0, K_\eta)$, where $K_\eta(\tau, u) = \sigma_\eta^2 \mathbf{1}(\tau = u)$ and $\mathbf{1}(\cdot)$ is the indicator function. We model each FLC ϕ_j as a smooth function admitting the basis expansion $\phi_j(\tau) = \boldsymbol{b}_\phi'(\tau)\boldsymbol{\xi}_j$, where $\boldsymbol{b}_\phi$ is a J_ϕ-dimensional vector of known basis functions and $\boldsymbol{\xi}_j$ is an unknown vector of coefficients. For superior MCMC performance, we prefer the low-rank thin plate spline basis for $\boldsymbol{b}_\phi$ (e.g., Crainiceanu et al., 2005) with knot locations selected using the quantiles of the observation points, $\mathcal{T}_o$. We place a smoothness prior on each $\boldsymbol{\xi}_j$, which is expressed

via a conditionally conjugate Gaussian distribution and is convenient for efficient posterior sampling (see the Appendix). The smoothness assumption typically produces more interpretable FLCs $\{\phi_j\}$ and can improve estimation for unobserved points $\tau^* \notin \mathcal{T}_o$. For the factors $\boldsymbol{e}_t = (e_{1,t}, \ldots, e_{J_\epsilon,t})'$, we assume $[\boldsymbol{e}_t|\boldsymbol{\Sigma}_e] \overset{indep}{\sim} N(\boldsymbol{0}, \boldsymbol{\Sigma}_e)$, with $\boldsymbol{\Sigma}_e = \text{diag}\left(\{\sigma_j^2\}_{j=1}^{J_\epsilon}\right)$ for simplicity. By comparison, the factors in Kowal et al. (2016) are time-dependent; we assume independence to obtain a special case of the FDLM in which the implied innovation process $\{\epsilon_t\}$ is an independent sequence, which also improves computational efficiency of the FDLM sampling algorithm. Importantly, we obtain a nonparametric, low-rank approximation to the innovation covariance, K_ϵ, with useful computational simplifications.

For identifiability, we order the factors according to variability of ϵ_t explained, $\sigma_1^2 > \sigma_2^2 > \cdots > \sigma_{J_\epsilon}^2 > 0$, and require orthonormality of the FLCs. It is computationally convenient to enforce the discrete orthonormality constraint $\boldsymbol{\Phi}'\boldsymbol{\Phi} = \boldsymbol{I}_{J_\epsilon}$, where $\boldsymbol{\Phi} = \boldsymbol{B}_\phi\boldsymbol{\Xi}$ is the $M \times J_\epsilon$ matrix of FLCs evaluated at $\mathcal{T}_e$, $\boldsymbol{B}_\phi = (\boldsymbol{b}_\phi(\tau_1), \ldots, \boldsymbol{b}_\phi(\tau_M))'$ is the $M \times J_\phi$ matrix of basis functions evaluated at $\mathcal{T}_e$, and $\boldsymbol{\Xi} = (\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_{J_\epsilon})$ is the $J_\phi \times J_\epsilon$ matrix of unknown FLC basis coefficients. The implied covariance matrix for $\boldsymbol{\epsilon}_t = (\epsilon_t(\tau_1), \ldots, \epsilon_t(\tau_M))'$ under (7) is $\boldsymbol{K}_\epsilon = \boldsymbol{\Phi}\boldsymbol{\Sigma}_e\boldsymbol{\Phi}' + \sigma_\eta^2\boldsymbol{I}_M$, conditional on $\{\phi_j, \sigma_j^2\}$ and σ_η^2. Importantly, the discretized orthonormality constraint offers a substantial simplification for computing the inverse of $\boldsymbol{K}_\epsilon$ using the Woodbury identity:

$$\boldsymbol{K}_\epsilon^{-1} = \sigma_\eta^{-2}\boldsymbol{I}_M - \sigma_\eta^{-2}\boldsymbol{\Phi}\tilde{\boldsymbol{\Sigma}}_e\boldsymbol{\Phi}', \tag{8}$$

where $\tilde{\boldsymbol{\Sigma}}_e = \sigma_\eta^{-2}\left(\boldsymbol{\Sigma}_e^{-1} + \sigma_\eta^{-2}\boldsymbol{\Phi}'\boldsymbol{\Phi}\right)^{-1} = \text{diag}\left(\{\sigma_j^2/(\sigma_\eta^2 + \sigma_j^2)\}_{j=1}^{J_\epsilon}\right)$. As a result, $\boldsymbol{K}_\epsilon^{-1}$ may be computed without any matrix inversions. By comparison, parametric covariance functions not only fail to offer computational simplifications for $\boldsymbol{K}_\epsilon^{-1}$, but also require *additional* computations of $\boldsymbol{K}_\epsilon^{-1}$ in the estimation of the covariance function parameters, $\boldsymbol{\rho}$. The FDLM sampling algorithm for the factors $\{e_{j,t}\}$, the FLCs $\{\phi_j\}$, and the variances $\{\sigma_j^2\}$ and σ_η^2 is computationally inexpensive and MCMC efficient. Note that the approximation error is a

nontrivial addition to model (7): η_t is necessary for nondegeneracy of $\boldsymbol{K}_\epsilon$, which is invertible only when $\sigma_\eta^2 > 0$. And while $\sigma_\eta^2 > 0$ implies that the innovations ϵ_t, and therefore μ_t, are not smooth, we find that in practice, the sample paths of ϵ_t and μ_t do appear smooth for sufficiently small σ_η^2. Generalizations to non-nugget approximation error variance functions $K_\eta(\tau, u) = \sigma_\eta^2(\tau)\mathbf{1}(\tau = u)$ for $\sigma_\eta^2 \colon \mathcal{T} \to \mathbb{R}^+$ are available, but may introduce additional model complexity and computational costs.

An important application of the FDLM simplification in (8) is given in Theorem 2, in which we derive a computationally convenient form for estimating the out-of-sample posterior distribution $[\mu_t(\tau^*)|\{\boldsymbol{y}_r\}_{r=1}^s]$ for $\tau^* \notin \mathcal{T}_e$, which includes as special cases the forecasting distribution ($s < t$), the filtering distribution ($s = t$), and the smoothing distribution ($s > t$).

4 Modeling the FAR Kernel

An accurate predictor of ψ is important not only for forecasting and inference, but also for interpreting the time evolution of $\{Y_t\}$. The likelihood for ψ is specified by the evolution equation in model (6), which may be extended for multiple lags. We select a Gaussian process prior for ψ, which encourages smoothness of the surface and produces more interpretable results. Using the basis approximation $\psi_\ell(\tau, u) = \boldsymbol{b}_0'(\tau, u)\boldsymbol{\theta}_{\psi_\ell}$, we place a Gaussian prior on $\boldsymbol{\theta}_{\psi_\ell}$, which induces a Gaussian process prior for ψ_ℓ. A tensor product basis $\boldsymbol{b}_0'(\tau, u) = (\boldsymbol{b}_\psi'(u) \otimes \boldsymbol{b}_\psi'(\tau))$ for $\boldsymbol{b}_\psi$ a J_ψ-dimensional vector of B-spline basis functions is computationally efficient in our setting, especially for large M. The details are presented in the Appendix. Since $J_\psi < M$, the evolution matrix $\boldsymbol{\Psi Q}$ in (6) has $J_\psi^2 < M^2$ unknown parameters, so the evolution equation in the DLM (6) has fewer parameters than a standard VAR(1) on $\boldsymbol{\mu}_t$. Notably, the posterior distribution for ψ_ℓ depends on $\boldsymbol{K}_\epsilon^{-1}$, which is computationally unstable for many common parametric covariance functions. By comparison, the nonparametric FDLM estimate of $\boldsymbol{K}_\epsilon^{-1}$ in (8) is computationally stable, which further stabilizes estimates of ψ_ℓ.

An important choice in the FAR(p) model is the maximum lag, p: a poor choice of p

can produce suboptimal forecasts and reduce MCMC efficiency. A reasonable approach is to compare the DIC or marginal likelihoods for different choices of p. However, this requires re-computing the model for each choice of p, which can be computationally intensive. Similarly, Kokoszka and Reimherr (2013) propose a multistage hypothesis testing procedure based on asymptotic approximations and an FPC decomposition, but would require modification for the hierarchical Bayesian implementation of (6).

Our approach is to select a maximum lag under consideration, p_{max}, and assign each lag ℓ a state variable, $s_\ell \in \{0, 1\}$, for $\ell = 1, \ldots, p_{max}$, to assess whether or not ψ_ℓ is included in the model:

$$\mu_t(\tau) = \sum_{\ell=1}^{p_{max}} s_\ell \int \psi_\ell(\tau, u)\mu_{t-\ell}(u)\, du + \epsilon_t(\tau), \tag{9}$$

which extends Kuo and Mallick (1998) and Korobilis (2013) to the FAR(p) setting. By averaging over the states $\{s_\ell\}_{\ell=1}^{p_{max}}$, the forecasts of model (9) are the model-averaged forecasts over the FAR(ℓ) models for $\ell = 1, \ldots, p_{max}$. Since we restrict $s_\ell \in \{0, 1\}$, rather than strongly shrinking ψ_ℓ toward zero, we can substantially improve computational efficiency: at each MCMC iteration, we sample $\{\boldsymbol{\mu}_t\}$ jointly from the FAR(p^*) extension of the DLM (6), where $p^* = \min\{\ell : s_{\ell+1} = \cdots = s_{p_{max}} = 0\}$ is the largest lag of nonzero autocorrelation.

The joint distribution of the states is $[s_1, s_2, \ldots, s_{p_{max}}] = [s_1] \prod_{\ell=2}^{p_{max}} [s_\ell | s_{\ell-1}, \ldots, s_1]$, where $[s_\ell | s_{\ell-1}, \ldots, s_1]$ is the probability that the lag ℓ autocorrelation term is included in the model, given whether the autocorrelation terms of the smaller lags $\ell - 1, \ldots, 1$ are included in the model. We assume that $s_\ell = 0$ implies that s_k is likely also zero for all $k > \ell$, which induces a more parsimonious model. In particular, we use the computationally convenient Markov assumption $[s_\ell | s_{\ell-1}, \ldots, s_1] = [s_\ell | s_{\ell-1}]$ with a small transition probability for $\mathbb{P}(s_\ell = 1 | s_{\ell-1} = 0) = q_{01}$. The reverse transition probability, $\mathbb{P}(s_\ell = 0 | s_{\ell-1} = 1) = q_{10}$, encourages smaller models when it is large. By default, we select $q_{01} = 0.01$, $q_{10} = 0.75$, and complete the joint prior distribution of $\{s_\ell\}_{\ell=1}^{p_{max}}$ with $\mathbb{P}(s_1 = 1) = 0.9$; for simulated data, the posterior does not appear to be sensitive to these choices.

5 Finite-Dimensional Optimality

The Gaussian assumptions in model (6) provide convenient posterior distributions for MCMC sampling and a useful framework for inference, but are not necessary for model (2). Suppose we relax the Gaussian assumption to $\epsilon_t \sim \mathcal{SP}(0, K_\epsilon)$, where $\mathcal{SP}(m, K)$ denotes a second-order stochastic process with mean function m and covariance function K. Similarly, let $\nu_{i,t}$ be a mean zero random variable with variance σ_ν^2 and let $\mu_1 \equiv \epsilon_1$. Given a finite set of evaluation points, $\mathcal{T}_e \subset \mathcal{T}$, model (4) implies the distribution-free DLM

$$\begin{cases} \boldsymbol{y}_t = \boldsymbol{Z}_t\boldsymbol{\mu} + \boldsymbol{Z}_t\boldsymbol{\mu}_t + \boldsymbol{\nu}_t, & \mathbb{E}[\boldsymbol{\nu}_t|\sigma_\nu^2] = 0,\ \mathrm{Cov}[\boldsymbol{\nu}_t|\sigma_\nu^2] = \sigma_\nu^2\boldsymbol{I}_{m_t}, \\ \boldsymbol{\mu}_t = \boldsymbol{\Psi}\boldsymbol{Q}\boldsymbol{\mu}_{t-1} + \boldsymbol{\epsilon}_t, & \mathbb{E}[\boldsymbol{\epsilon}_t|\boldsymbol{K}_\epsilon] = 0,\ \mathrm{Cov}[\boldsymbol{\epsilon}_t|\boldsymbol{K}_\epsilon] = \boldsymbol{K}_\epsilon, \end{cases} \tag{10}$$

under the integral approximation (5), where the vectors and matrices are defined as before and $\boldsymbol{\mu}_1 \equiv \boldsymbol{\epsilon}_1$. Since this holds for any finite set of evaluation points $\mathcal{T}_e \subset \mathcal{T}$, we may consider the DLM (10) to be a *collection* of models indexed by the evaluation points, $\mathcal{T}_e$. The error sequences, $\{\boldsymbol{\nu}_t\}$ and $\{\boldsymbol{\epsilon}_t\}$, are assumed to be uncorrelated, rather than independent. If we additionally assume Gaussianity of $\{\boldsymbol{\nu}_t\}$ and $\{\boldsymbol{\epsilon}_t\}$, then the uncorrelatedness implies independence, and model (10) becomes model (6). Extensions for the FAR(p) models are similar. The results below also hold for time-dependent variances for $\boldsymbol{\nu}_t$ and $\boldsymbol{\epsilon}_t$.

Let d be an estimator of $\delta \in L^2(\mathcal{T})$, and consider the squared error loss using the Euclidean norm: $\mathcal{L}_e(\delta, d) = (\boldsymbol{\delta} - \boldsymbol{d})'(\boldsymbol{\delta} - \boldsymbol{d})$, where $\mathcal{L}_e$ is indexed by the set of evaluation points, $\mathcal{T}_e$, at which δ and d are evaluated to form the corresponding vectors $\boldsymbol{\delta}$ and $\boldsymbol{d}$. When $\mathcal{T}_e$ is an equally-spaced fine grid on $\mathcal{T}$, the loss function $\mathcal{L}_e$ will approximate the usual loss function for functional data, $\mathcal{L}_{L^2}(\delta, d) = \int (\delta(u) - d(u))^2\, du$, for most reasonable choices of δ and d (up to a rescaling by $M = |\mathcal{T}_e|$). In a standard Bayesian analysis, the goal would be to minimize the posterior risk, $\mathbb{E}[\mathcal{L}_e(\delta, d)|\{\boldsymbol{y}_t\}]$, for which the solution is the posterior expectation, $d = \mathbb{E}[\delta|\{\boldsymbol{y}_t\}]$. Indeed, the estimators discussed below minimize the posterior risk under the Gaussian assumptions of model (6). However, by relaxing the distributional

assumptions in (10) to increase the generality of the model, we no longer have sufficient information to compute posterior distributions or posterior moments. In addition, it is difficult to compare Bayesian and non-Bayesian procedures under the posterior risk, and most procedures for functional time series modeling are non-Bayesian. Therefore, we consider the overall risk $\mathcal{R}_e(\delta, d) = \mathbb{E}[\mathcal{L}_e(\delta, d)]$, which is the expected value of the posterior risk with respect to the sampling distribution. As with the loss function $\mathcal{L}_e$, the risk function $\mathcal{R}_e$ is indexed by the evaluation points, $\mathcal{T}_e$; we seek to minimize $\mathcal{R}_e$ for any choice of $\mathcal{T}_e$.

Let $\mathcal{D}_t = \{\boldsymbol{y}_t, \boldsymbol{y}_{t-1}, \ldots, \boldsymbol{y}_1\} \cup \mathcal{D}_0$ be the information available at time t, where $\mathcal{D}_0$ represents the information prior to time $t = 1$.

Theorem 1. *For* any *finite set of evaluation points* $\mathcal{T}_e \subset \mathcal{T}$*, the* unique best linear predictor *of the conditional random vector* $\boldsymbol{\delta} \sim [\boldsymbol{\delta}|\boldsymbol{Y}, \boldsymbol{\Theta}]$*, where* $\boldsymbol{\delta}, \boldsymbol{Y} \subseteq \mathcal{D}_T \cup \{\mu_t(\tau) : \tau \in \mathcal{T}_e, t = 1, \ldots, T\}$ *and* $\boldsymbol{\Theta} = \{\mu, \sigma_\nu^2, \psi, K_\epsilon\}$*, under the risk* $\mathcal{R}_e$ *and conditional on model* (4) *with the integral approximation* (5)*, is the conditional expectation* $\hat{\boldsymbol{\delta}}(\boldsymbol{Y}|\boldsymbol{\Theta}) \equiv \mathbb{E}[\boldsymbol{\delta}|\boldsymbol{Y}, \boldsymbol{\Theta}]$ *as computed under model* (6)*.*

The proof of Theorem 1 is in the Appendix, and extends fundamental results for vector-valued DLMs. The best linear predictors of Theorem 1 equivalently minimize the risk $\mathcal{R}(\delta, d) = \sup_{\mathcal{T}_e} \mathcal{R}_e(\delta, d)$ among all linear estimators, where the sup is taken over all finite $\mathcal{T}_e \subset \mathcal{T}$. The most useful examples of $[\boldsymbol{\delta}|\boldsymbol{Y}, \boldsymbol{\Theta}]$ in Theorem 1 are the forecasting distributions $[\boldsymbol{y}_{t+h}|\mathcal{D}_t, \boldsymbol{\Theta}]$ and $[\boldsymbol{\mu}_{t+h}|\mathcal{D}_t, \boldsymbol{\Theta}]$ for $h > 0$, the smoothing distributions $[\boldsymbol{\mu}_t|\mathcal{D}_T, \boldsymbol{\Theta}]$, and the filtering distributions $[\boldsymbol{\mu}_t|\mathcal{D}_t, \boldsymbol{\Theta}]$, for $t = 1, \ldots, T$. Theorem 1 depends on the observation points $\mathcal{T}_o$ only via the assumption that $\boldsymbol{Z}_t$ is known. In general, we assume $\mathcal{T}_o \subseteq \mathcal{T}_e$, so $\boldsymbol{Z}_t$ is an incidence matrix and therefore known. Theorem 1 does *not* require $\mathcal{T}_o$ to become arbitrarily dense in $\mathcal{T}$, and is valid for both sparse and dense designs. For implementation, we compute the relevant expectations within the Gibbs sampling algorithm (see Section A of the web supplement), and then average over the Gibbs sample of $\boldsymbol{\Theta}$. Alternatively, an EM algorithm could be used to estimate the relevant expectations (Cressie and Wikle, 2011).

There is no intrinsic reason to restrict the estimators to linearity. However, several popular competing methods are linear, and therefore are dominated by the conditional expectations computed from model (6) whenever the estimators are distinct. More formally:

Corollary 1. *Consider a basis expansion of the observations* $\boldsymbol{y}_t \approx \boldsymbol{B}_t\boldsymbol{\theta}_t$*, where* $\boldsymbol{B}_t' = (\boldsymbol{b}(\tau_{1,t}), \ldots, \boldsymbol{b}(\tau_{m_t,t}))$*,* $\boldsymbol{b}$ *is a known* J*-dimensional vector of basis functions, and* $\boldsymbol{\theta}_t$ *is the corresponding* J*-dimensional vector of unknown basis coefficients. If the estimator* $\hat{\boldsymbol{\theta}}_t$ *of* $\boldsymbol{\theta}_t$ *is linear in* $\boldsymbol{y}_t$*, then estimates or forecasts of the form* $\boldsymbol{H}\hat{\boldsymbol{\theta}}_t + \boldsymbol{h}$*, conditional on the matrix* $\boldsymbol{H}$ *and the vector* $\boldsymbol{h}$*, are inadmissible for all* $[\boldsymbol{\delta}|\boldsymbol{Y}]$ *whenever* $\boldsymbol{H}\hat{\boldsymbol{\theta}}_t + \boldsymbol{h} \neq \hat{\boldsymbol{\delta}}(\boldsymbol{Y}|\boldsymbol{\Theta})$*.*

The most important application of Corollary 1 is to characterize the inadmissibility of procedures based on FPC scores. In the notation of Corollary 1, let $\boldsymbol{b}$ be the FPC basis, which we assume is fixed and known. The components of $\boldsymbol{\theta}_t$ correspond to the FPC scores, defined by $\theta_{j,t} = \int \{Y_t(u) - \mu(u)\} b_j(u)\, du = \int \mu_t(u) b_j(u)\, du$. There are two standard approaches for computing FPC scores: quadrature methods for dense designs absent measurement error, and the PACE procedure of Yao et al. (2005), which uses conditional expectations under a Gaussian assumption and applies more generally. In both cases, the FPC scores are linear in $\boldsymbol{y}_t$, so Corollary 1 applies.

Among functional time series methods, the most pertinent procedures are Aue et al. (2015) and Hyndman and Ullah (2007). Aue et al. (2015) provide the more general framework, in which they compute the best linear predictors for the FPC scores, and then forecast the FPC scores using multivariate time series methods. For time series methods that are *linear* in the FPC scores, such forecasts are inadmissible. While Aue et al. (2015) undoubtedly provide a simple yet general framework for forecasting a functional time series, the simulations of Section 6 confirm the consequence of inadmissibility on forecasting performance.

Corollary 2. *Consider the common functional data pre-processing procedure in which the discrete, noisy observations,* $\boldsymbol{y}_t$*, are replaced by estimated functions evaluated on a fine grid,* $\hat{\boldsymbol{y}}_t$*, and then estimates and forecasts are computed using the functional "data"* $\hat{\boldsymbol{y}}_t$*. If* $\hat{\boldsymbol{y}}_t$ *is*

linear in $\{\boldsymbol{y}_t\}$*, then any estimator or forecast linear in* $\{\hat{\boldsymbol{y}}_t\}$ *is inadmissible for all* $[\boldsymbol{\delta}|\boldsymbol{Y}]$ *whenever* $\hat{\boldsymbol{y}}_t \neq \hat{\boldsymbol{\delta}}(\boldsymbol{Y}|\boldsymbol{\Theta})$.

Typically, $\hat{\boldsymbol{y}}_t$ is estimated using splines or kernel smoothers, both of which are linear in $\boldsymbol{y}_t$. As an application of Corollary 2, the simple forecasting method of fitting a VAR to $\hat{\boldsymbol{y}}_t$ evaluated on a grid of points, conditional on the VAR coefficient matrix, is inadmissible.

Corollary 3. *The unique best linear predictor of* $[\mu_t(\tau^*)|\mathcal{D}_s]$ *for any times* t, s *and any point* $\tau^* \in \mathcal{T}$ *is the corresponding expectation under model* (6).

Model (6) achieves the optimality of a kriging estimator for interpolation of any point $\tau^* \in \mathcal{T}$, simply by adding τ^* to the evaluation set $\mathcal{T}_e$.

In practice, we need not include all such τ^* in $\mathcal{T}_e$: we can estimate the out-of-sample posterior distribution $[\mu_t(\tau^*)|\mathcal{D}_s]$ for $\tau^* \notin \mathcal{T}_e$ by sampling from the out-of-sample full conditional distribution $\left[\mu_t(\tau^*)|\{\boldsymbol{\mu}_r\}_{r=1}^T, \boldsymbol{\Theta}, \mathcal{D}_s\right]$ within the Gibbs sampler, and then averaging over the Gibbs sample of $\{\boldsymbol{\mu}_r\}_{r=1}^T$ and $\boldsymbol{\Theta}$. Let $\boldsymbol{\psi}'(\tau^*) \equiv (\psi(\tau^*, \tau_1), \ldots, \psi(\tau^*, \tau_M))$ and $\boldsymbol{\phi}'(\tau^*) \equiv (\phi_1(\tau^*), \ldots, \phi_{J_\epsilon}(\tau^*))$. In the special case of model (4) and using the FDLM (7), we have the following computationally efficient alternative for state space imputation:

Theorem 2. *Suppose* $\tau^* \in \mathcal{T}$ *such that* $\tau^* \notin \mathcal{T}_e$*. Under the FDLM* (7) *and conditional on model* (4) *with the integral approximation* (5)*, the out-of-sample full conditional distribution of* $\mu_t(\tau^*)$ *is* $\left[\mu_t(\tau^*)|\{\boldsymbol{\mu}_r\}_{r=1}^T, \boldsymbol{\Theta}, \mathcal{D}_s\right] \sim N\left(m_t(\tau^*), K_t(\tau^*)\right)$*, where* $m_t(\tau^*) = \boldsymbol{\psi}'(\tau^*)\boldsymbol{Q}\boldsymbol{\mu}_{t-1} + \boldsymbol{\phi}'(\tau^*)\tilde{\boldsymbol{\Sigma}}_e\boldsymbol{\Phi}'\left(\boldsymbol{\mu}_t - \boldsymbol{\Psi}\boldsymbol{Q}\boldsymbol{\mu}_{t-1}\right)$ *and* $K_t(\tau^*) = \sigma_\eta^2 + \sigma_\eta^2\boldsymbol{\phi}'(\tau^*)\tilde{\boldsymbol{\Sigma}}_e\boldsymbol{\phi}(\tau^*)$.

The proof of Theorem 2 and extensions for $p > 1$ are in Section B.3 of the web supplement. Using Theorem 2, we can efficiently estimate the out-of-sample posterior distribution $[\mu_t(\tau^*)|\mathcal{D}_s]$ with minimal adjustments to the Gibbs sampling algorithm (see Section A.2 of the web supplement). Theorem 2 builds upon the approximation in (5) and the computational simplifications of the FDLM to produce simple and efficient moment calculations for the full conditional distributions without expanding the dimension of the state vector,

M. Note that for implementation purposes, the terms $\boldsymbol{\mu}_t$ and $\boldsymbol{\mu}_{t-1}$ appearing in $m_t(\tau^*)$ are assumed to be sampled from the full conditional distribution $\left[\{\boldsymbol{\mu}_r\}_{r=1}^T | \boldsymbol{\Theta}, \mathcal{D}_s\right]$.

6 Simulations

We conducted extensive simulations to evaluate the proposed methods for FAR(p) relative to several competitive alternatives. We are particularly interested in one-step forecasting and recovery of the FAR kernel ψ_1, and in how the associated performance varies with the sample size T, the location and number of the observation points $\tau_{1,t}, \ldots, \tau_{m_t,t}$, the kernel ψ_1, and the smoothness of the innovation process ϵ_t. We also assess the performance of the model averaging procedure of Section 4 for $p \in \{1, 2\}$, and compare the nonparametric FDLM approach of Section 3 with a more standard parametric Gaussian process implementation.

6.1 Sampling Designs

For all simulations, the mean function is $\mu(\tau) = \frac{1}{10}\tau^3 \sin(2\pi\tau)$, which produces the dominate shape in the rightmost panels of Figure 1. The measurement errors are identically distributed for all simulations: $\nu_{i,t} \overset{iid}{\sim} N(0, \sigma_\nu^2)$ with $\sigma_\nu = 0.002$. We vary the sample size from small ($T = 50$) to large ($T = 350$) for the FAR(1) simulations, and use a moderate sample size ($T = 125$) for the FAR(2) simulation. The FAR(1) kernel used for Figure 1 is the *Bimodal-Gaussian* kernel, $\psi(\tau, u) \propto \frac{0.75}{\pi(0.3)(0.4)} \exp\{-(\tau - 0.2)^2/(0.3)^2 - (u - 0.3)^2/(0.4)^2\} + \frac{0.45}{\pi(0.3)(0.4)} \exp\{-(\tau - 0.7)^2/(0.3)^2 - (u - 0.8)^2/(0.4)^2\}$, following Wood (2003); see the web supplement for a plot of the Bimodal-Gaussian kernel. We also present results for the *Linear-τ* kernel, $\psi(\tau, u) \propto \tau$, and the *Linear-$u$* kernel, $\psi(\tau, u) \propto u$. Each kernel is rescaled according to a pre-specified squared norm, $C_{\psi_\ell} = \int\int \psi_\ell^2(\tau, u)\, d\tau\, du$, with $\sum_{\ell=1}^p C_{\psi_\ell} < 1$ for stationarity. We select $C_{\psi_1} = 0.8$ for the FAR(1) simulations and use $(C_{\psi_1}, C_{\psi_2}) = (0.4, 0.2)$ for the FAR(2) simulation; smaller values of C_{ψ_ℓ} produce similar comparative results, but the forecasting performance deteriorates for all methods. For the

innovation process, ϵ_t, we consider both smooth and non-smooth Gaussian processes. We use the covariance function parametrization $K_\epsilon = \sigma^2 R_\rho$, where R_ρ is the Matérn correlation function $R_\rho(\tau, u) = \{2^{\rho_1 - 1}\Gamma(\rho_1)\}^{-1} \left(||\tau - u||/\rho_2\right)^{\rho_1} K_{\rho_1}(||\tau - u||/\rho_2)$, $\Gamma(\cdot)$ is the gamma function, K_{ρ_1} is the modified Bessel function of order ρ_1, and $\boldsymbol{\rho} = (\rho_1, \rho_2)$ are parameters (Matérn, 2013). We let $\sigma = 0.01$ and $\boldsymbol{\rho} = (\rho_1, 0.1)$, with $\rho_1 = 2.5$ for smooth (twice-differentiable) sample paths and $\rho_1 = 0.5$ for non-smooth (continuous, non-differentiable) sample paths.

We consider three sampling designs for the observation points: *dense*, *sparse-random*, and *sparse-fixed*. In each case, the set of evaluation points, $\mathcal{T}_e$, is an equally-spaced grid of $M = 30$ points on $\mathcal{T} = [0, 1]$. The *dense* design uses $m_t = 25$ equally-spaced observation points on $[0, 1]$ for all t, for which the results are representative of denser ($m_t \gg 25$) designs and similar to those of Didericksen et al. (2012); see Section C of the web supplement. The *sparse-random* design is generated by first sampling each m_t from a zero-truncated Poisson (5) distribution, and then sampling $\tau_{1,t}, \ldots, \tau_{m_t,t}$ without replacement from $\mathcal{T}_e$. This is a common design in sparse functional data, in which m_t may be small for some t, but $\mathcal{T}_o$ is dense in $\mathcal{T}$. The *sparse-fixed* design uses $m_t = 8$ equally-spaced points in $\mathcal{T}$. This is the most challenging design, and one for which multivariate time series methods should be most competitive with functional time series methods. Comparatively, the sparse settings are similar to the dense setting, but with additional missing observations.

6.2 Competing Estimators

Within the proposed framework and using the FDLM of Section 3 for the innovation covariance function, we compute forecasts for $p = 1$ (FDLM-FAR(1)), and in the FAR(2) simulation, for $p = 2$ (FDLM-FAR(2)) and $p = 3$ (FDLM-FAR(3)). We also compute forecasts using the model averaging procedure with $p_{max} = 4$ (FDLM-FAR(p)). To assess the performance of the FDLM implementation, we compute forecasts using model (6) with a parametric covariance function for $K_\epsilon = \sigma^2 R_\rho$ (GP-FAR(1)). We use the Matérn correlation function for R_ρ, with $\rho_1 = 2.5$ as in the smooth Gaussian process simulations, and use the

priors $\sigma^{-2} \sim \text{Gamma}\,(10^{-3}, 10^{-3})$ and $\rho_2 \sim \text{Uniform}\,(0, U_{\rho_2})$, where U_{ρ_2} is the maximum value of ρ_2 for which the correlation function R_ρ is less than 0.99 for all pairs of evaluation points. These models are implemented using the Gibbs sampling algorithm provided in Section A of the web supplement, and estimates are based on 5,000 MCMC simulations after a burn-in of 5,000. For the large sample setting ($T = 350$), the mean computation time per 1,000 MCMC simulations was 2.3 minutes for FDLM-FAR(1) and 4.4 minutes for GP-FAR(1). The computing times are calculated on a 64-bit Windows machine with a 2.40-GHz Intel core i7-4700MQ processor with 8 GB of RAM, and the code is written in `R`.

We consider several important competing methods. Let $\hat{\boldsymbol{y}}_{t+1}$ denote the one-step forecast at time t. For baseline comparisons, we use the random-walk (RW) forecast, $\hat{\boldsymbol{y}}_{t+1} = \boldsymbol{y}_t$, and the mean (Mean) forecast, $\hat{\boldsymbol{y}}_{t+1} = \hat{\boldsymbol{\mu}}$, where $\hat{\boldsymbol{\mu}}$ is a smooth estimate of the mean of $\{\boldsymbol{y}_s\}_{s=1}^t$. We estimate $\hat{\boldsymbol{\mu}}$ using a B-spline basis expansion via the function `meanfd()` in the `R` package `fda` (Ramsay et al., 2014). Both estimators are robust against overfitting, and the mean forecast is optimal when $\psi = 0$. We also compute the one-step forecast based on a VAR(1) fit to $\{\boldsymbol{y}_s\}_{s=1}^t$ (VAR-Y). In the sparse-random design, the observations $\boldsymbol{y}_t$ were used to linear interpolate on $\mathcal{T}_e$ prior to fitting the VAR. In the sparse-fixed design, the VAR was fit to the observation points, and then forecasts for the evaluation points were computed by fitting a spline to the VAR forecasts of the observation points. For additional comparisons, we computed forecasts from a simple exponential smoother (SES) applied pointwise to each component of $\boldsymbol{y}_t$, i.e., each time series $\{y_{j,t}\}_{t=1}^T$. The SES forecasts are implemented using the `ses` function in the `R` package `forecast` (Hyndman and Khandakar, 2008), with an identical imputation scheme as VAR-Y. We also considered two functional data methods. First, we used the Estimated Kernel procedure outlined in Horváth and Kokoszka (2012), which estimates ψ_ℓ in (2) using FPCs (FAR Classic); we fix $p = 1$ for simplicity. This method has well-studied theoretical properties and is a useful baseline for FAR models. Second, we implemented the method of Aue et al. (2015), which we briefly described in Section 5, using a VAR(1) on the FPC scores (VAR-FPC). We compute the FPCs using

the `fda` package in `R` with B-spline basis functions. To avoid the ill-conditioned estimators discussed in Horváth and Kokoszka (2012), we regularize via basis truncation, using 8 equally-spaced interior knots. The number of components is selected to explain at least 95% of the variability in $\{\boldsymbol{y}_t\}$. For the sampling designs considered here, this approach works well. Finally, we report the oracle forecast (FAR Oracle) computed using the true one-step forecasts $\mathbb{E}\left[\mu_t(\tau)|\{\psi_\ell, \mu_{t-\ell}\}_{\ell=1}^p\right] = \sum_{\ell=1}^p \int \psi_\ell(\tau, u)\mu_{t-\ell}(u)\, du$ within the simulation, where $\{\psi_\ell\}_{\ell=1}^p$ are the FAR kernels from the simulation specification, $\{\mu_t\}$ are the simulated values of the latent FAR process, and the integral is approximated using the trapezoidal rule with $M = 200$ grid points. The oracle forecast is not actually an estimator, and is unaffected by sparsity or small sample sizes.

We estimate the one-step forecasts $[\boldsymbol{y}_{T+h}|\boldsymbol{y}_{1:(T+h-1)}]$, $h = 1, \ldots, 25$, for all estimators under consideration, and compare them using the mean squared forecast error $MSFE_e = \frac{1}{25M}\sum_{h=1}^{25} ||\boldsymbol{Y}_{T+h} - \hat{\boldsymbol{Y}}_{T+h}||^2$ where $\boldsymbol{Y}_{T+h} = (Y_{T+h}(\tau_1), \ldots, Y_{T+h}(\tau_M)))'$, which measures the one-step forecasting performance at the *evaluation* points, and the mean squared error $MSE_{\psi_1} = \frac{1}{M^2}\sum_{i=1}^M \sum_{k=1}^M \{\psi_1(\tau_i, \tau_k) - \hat{\psi}_1(\tau_i, \tau_k)\}^2$, which measures the recovery of the lag-1 kernel ψ_1. Because $\mathcal{T}_e$ is relatively dense in $\mathcal{T}$, $MSFE_e$ and MSE_{ψ_1} approximate the integrated squared errors $\int \{Y_{T+h}(u) - \hat{Y}_{T+h}(u)\}^2\, du$ and $\int\int \{\psi_1(\tau, u) - \hat{\psi}_1(\tau, u)\}^2\, d\tau\, du$, respectively. Estimators $\hat{\psi}_1$ are available only for the proposed methods and FAR Classic. For computational convenience in the proposed methods, we update $\{\boldsymbol{\mu}_t\}_{t=1}^{T+h-1}$ using all of the data $\boldsymbol{y}_{1:(T+h-1)}$, but sample all other parameters only conditional on $\boldsymbol{y}_{1:T}$. DLM updating algorithms provide recursive one-step forecasts for $\boldsymbol{\mu}_t$, but in general there are no convenient updating algorithms for the other parameters. In practice, this is not a problem, but suggests that our simulation analysis may underestimate the performance of the proposed model.

6.3 Results

We computed $MSFE_e$ and MSE_{ψ_1} under a variety of sampling designs, each for $N = 50$ simulations, and present the results for a few important cases in Figures 2 and 3, respec-

tively. The figures are color-coded: multivariate methods are green, existing functional data methods are red, the proposed methods are blue, and the oracle is gold.

For the sparse designs in Figure 2, the proposed methods are all superior to the competitors, and in some cases nearly achieve the oracle performance, even though the oracle is unaffected by sparsity. Figure 3 shows that the proposed methods also offer a substantial improvement in ψ_1 estimation. Importantly, the proposed model with model averaging is competitive with the known p model for both forecasting and estimation of ψ_1. The model averaging procedure of Section 4 typically identifies the true p with high probability, with a mild tendency to overestimate p. However, this behavior is encouraging: the bottom right panel of Figure 3, in which $p = 2$, suggests that overestimating the lag (FDLM-FAR(3)) is preferable to underestimating the lag (FDLM-FAR(1), GP-FAR(1)) for ψ_1 estimation. FDLM-FAR(1) is competitive with GP-FAR(1), even when the parametric Gaussian process model assumes the correct (smooth) innovation distribution, which suggests that the FDLM implementation of Section 3 provides an adequate approximation. Under the dense design (see Section C of the web supplement), the improvements of the proposed methods over existing functional data methods are less substantial, and for $T = 350$ the functional data methods all nearly achieve the oracle performance. The proposed methods, however, again provide superior recovery of ψ_1. In general, we find that the functional data methods, in particular the proposed approaches, outperform the multivariate methods, especially in the dense design. We conclude that the proposed methods provide highly competitive forecasts and superior FAR kernel recovery in a wide variety of important settings.

7 Forecasting Nominal and Real Yield Curves

We apply the proposed methods to model and forecast nominal and real yield curves. Yield curves are important in a variety of economic and financial applications, such as evaluating economic and monetary conditions, pricing fixed-income securities, generating forward

curves, computing inflation premiums, and monitoring business cycles (Bolder et al., 2004). In practice, the U.S. real yield curve is estimated using Treasury Inflation-Protected Securities (TIPS), for which payments are adjusted according to the Consumer Price Index for All Urban Consumers (CPI-U) to provide investors with protection against inflation. U.S. nominal and TIPS yield curve data are published daily by the Federal Reserve, which uses actively-traded securities to fit a quasi-cubic spline for each curve. Estimates of the real and nominal yield curves are provided for maturities $\mathcal{T}_t^R = \{60, 84, 120, 240, 360\}$ and $\mathcal{T}_t^N = \{1, 3, 6, 12, 24, 36\} \cup \mathcal{T}_t^R$ months, respectively. Notably, the real yield is observed sparsely, and only at longer maturities. The small number of available maturities for real yields presents a challenge for existing functional time series models, and provides an interesting comparison with the nominal yield, for which there are more observed maturities.

To assess the performance of the proposed model, we conducted an extensive forecasting study using daily nominal and real yield curve data. Beginning in 2003, we construct nine consecutive yet non-overlapping 18-month subperiods for estimation ($T \approx 375$); the corresponding starting dates are given in Table 1. For the month following each estimation period, we compute both one- and five-step (i.e., one business week) forecasts (≈ 20 and ≈ 15 time points, respectively) for *both* the nominal and real yields. In all cases, the nominal and real yields are modeled separately in order to provide additional comparisons.

We compute forecasts for the proposed methods by simulating from the forecasting distribution in the DLM (6). For computational convenience, we update only the DLM state parameters $\{\boldsymbol{\mu}_t\}$ during the forecast periods, and fix the remaining parameters based on the estimation periods. We also rescale the observation points $\mathcal{T}_t^R$ and $\mathcal{T}_t^N$ such that $\mathcal{T}_t^R, \mathcal{T}_t^N \subset \mathcal{T} = [0, 1]$. We compute forecasts using the competing methods described in Section 6, which use all available data for each forecast. For further comparisons, we include two popular parametric yield curve models based on the Nelson-Siegel parametrization (Nelson and Siegel, 1987): Diebold and Li (2006, DL), which extends the Nelson-Siegel model to the dynamic setting via a two-step estimation procedure, and Diebold et al. (2006, DRA)

which is similar to DL, but instead estimates parameters jointly using maximum likelihood within a state space model; see the web supplement for implementation details.

The one- and five-step root mean squared forecasting errors (RMSFEs) for the nominal yields and real yields are in Tables 1 and 2, respectively. We omit unstable DRA forecasts, as well as multi-step forecasts for FAR Classic, which are unavailable. For both data sets, the proposed methods—denoted FAR(1) and FAR(p), using the lag selection procedure with $p_{max} = 3$—are consistently among the best forecasters for all time periods, and outperform the existing functional data forecasts by a wide margin. For the nominal yields, the FAR(1) provides the best one-step forecasts aggregated across all time periods. For the real yields, the proposed methods are again among the most competitive, particularly in the periods since the financial crisis. Echoing the results in Diebold and Li (2006), the RW forecast is a difficult benchmark to clear, and the existing functional data models typically fail to do so. By comparison, the proposed FAR forecasts are highly competitive across all time periods and for both the nominal and (sparsely-observed) real yields.

An important feature of the proposed FAR model is the ability to compute exact (up to MCMC error) credible bands for parameters of interest, including forecasts. Such uncertainty quantification is unavailable for the RW forecast, which is our primary competitor in this application. For illustration, we compute pointwise and simultaneous credible bands for one-step forecasts during August 2016 in Figure 4. For both nominal and real yields, the credible bands are tighter for shorter maturities and widen in regions of unobserved points, which is appropriate behavior for a nonparametric method.

8 Concluding Remarks

The proposed hierarchical FAR(p) model provides a useful framework for estimation, inference, and forecasting functional time series data. Our model is especially suited for sparsely or irregularly sampled curves and for curves sampled with non-negligible measurement error,

and produces best linear predictors in a general FAR(p) setting, thereby dominating many competing functional time series models. The FDLM provides a more flexible, computationally efficient, and stable approach for modeling (innovation) covariance functions. Our model averaging procedure provides an effective solution to the problem of specifying p, and produces highly competitive forecasts. The simulation analysis and yield curve application suggest that the proposed FAR(p) model may improve forecasting and estimation in a wide range of settings, and the efficient MCMC sampling algorithm allows us to perform exact (up to MCMC error and prior misspecification) inference for important parameters.

While we assumed independent factors (and therefore independent innovations) in Section 3, we can relax this assumption and allow $\mathbf{\Sigma}_e$ to be a stochastic process evolving over time. In this more general framework, the FDLM (7) can accommodate stochastic volatility or heavier-tailed distributions for the factors, yet retains the computational simplifications of (8) and Theorem 2. Letting $\mathbf{\Sigma}_t = \text{diag}\left(\{\sigma^2_{j,t}\}_{j=1}^{J_\epsilon}\right)$, the (time-dependent) innovation covariance function is $K_{\epsilon_t}(\tau, u) \equiv \text{Cov}\left(\epsilon_t(\tau), \epsilon_t(u)\right) = \sum_{j=1}^{J_\epsilon} \sigma^2_{j,t}\phi_j(\tau)\phi_j(u) + \sigma^2_\eta \mathbf{1}\{\tau = u\}$. By modeling each $\{\sigma^2_{j,t}\}_{t=1}^T$ for $j = 1, \ldots, J_\epsilon$ with an independent stochastic volatility model (e.g., Kim et al., 1998), the time-dependence of $\{\sigma^2_{j,t}\}$ will propagate to the innovation covariance functions, K_{ϵ_t}. Similar modifications can accommodate scale-mixtures of Gaussian distributions for the factors (Fernandez and Steel, 2000) to induce more general distributions for the innovation process, $\{\epsilon_t\}$. These generalizations are particularly important for financial applications, for which stochastic volatility models and heavy-tailed distributions are commonly appropriate.

Future work will investigate more adaptive FAR(p) models for longer, possibly nonstationary functional time series through stochastic volatility, time-varying ψ_ℓ, and regime shifts. Important extensions also include modeling multiple functional responses $Y_t(\tau) \in \mathbb{R}^d$ for $d > 1$, which requires a model for both the auto- and cross-correlations, and incorporating exogenous predictors. In both cases, the DLM framework of (6) offers a promising platform for pursuing these extensions.

References

Aue, A., Norinho, D. D., and Hörmann, S. (2015). On the prediction of stationary functional time series. *Journal of the American Statistical Association*, 110(509):378–392.

Behseta, S., Kass, R. E., and Wallstrom, G. L. (2005). Hierarchical models for assessing variability among functions. *Biometrika*, 92(2):419–434.

Besse, P. C., Cardot, H., and Stephenson, D. B. (2000). Autoregressive forecasting of some functional climatic variations. *Scandinavian Journal of Statistics*, pages 673–687.

Bolder, D., Johnson, G., and Metzler, A. (2004). *An empirical analysis of the Canadian term structure of zero-coupon interest rates*. Bank of Canada.

Bosq, D. (2000). *Linear processes in function spaces: theory and applications*, volume 149. Springer Science & Business Media.

Bosq, D. and Blanke, D. (2008). *Inference and prediction in large dimensions*, volume 754. John Wiley & Sons.

Cardot, H., Ferraty, F., and Sarda, P. (1999). Functional linear model. *Statistics & Probability Letters*, 45(1):11–22.

Chen, Y. and Li, B. (2015). An adaptive functional autoregressive forecast model to predict electricity price curves. *Journal of Business & Economic Statistics*, (just-accepted):1–56.

Crainiceanu, C., Ruppert, D., and Wand, M. P. (2005). Bayesian analysis for penalized spline regression using WinBUGS. *Journal of Statistical Software*, 14(14):1–24.

Cressie, N. and Wikle, C. K. (2011). *Statistics for spatio-temporal data*. John Wiley & Sons.

Damon, J. and Guillas, S. (2002). The inclusion of exogenous variables in functional autoregressive ozone forecasting. *Environmetrics*, 13:759–774.

Damon, J. and Guillas, S. (2005). Estimation and simulation of autoregressive hilbertian processes with exogenous variables. *Statistical Inference for Stochastic Processes*, 8(2):185–204.

Didericksen, D., Kokoszka, P., and Zhang, X. (2012). Empirical properties of forecasts with the functional autoregressive model. *Computational Statistics*, 27(2):285–298.

Diebold, F. X. and Li, C. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics*, 130(2):337–364.

Diebold, F. X., Rudebusch, G. D., and Aruoba, B. S. (2006). The macroeconomy and the yield curve: a dynamic latent factor approach. *Journal of Econometrics*, 131(1):309–338.

Durbin, J. and Koopman, S. J. (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika*, 89(3):603–616.

Earls, C. and Hooker, G. (2014). Bayesian covariance estimation and inference in latent Gaussian process models. *Statistical Methodology*, 18:79–100.

Fernandez, C. and Steel, M. F. (2000). Bayesian regression analysis with scale mixtures of normals. *Econometric Theory*, 16(01):80–101.

Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Analysis*, 1(3):515–534.

Hays, S., Shen, H., and Huang, J. Z. (2012). Functional dynamic factor models with application to yield curve forecasting. *The Annals of Applied Statistics*, 6(3):870–894.

Horváth, L. and Kokoszka, P. (2012). *Inference for functional data with applications*, volume 200. Springer Science & Business Media.

Hyndman, R. and Khandakar, Y. (2008). Automatic time series forecasting: The forecast package for R. *Journal of Statistical Software*, 27(1):1–22.

Hyndman, R. J. and Ullah, M. S. (2007). Robust forecasting of mortality and fertility rates: a functional data approach. *Computational Statistics & Data Analysis*, 51(10):4942–4956.

Kargin, V. and Onatski, A. (2008). Curve forecasting by functional autoregression. *Journal of Multivariate Analysis*, 99(10):2508–2526.

Kaufman, C. G., Sain, S. R., et al. (2010). Bayesian functional ANOVA modeling using gaussian process prior distributions. *Bayesian Analysis*, 5(1):123–149.

Kim, S., Shephard, N., and Chib, S. (1998). Stochastic volatility: likelihood inference and comparison with ARCH models. *The Review of Economic Studies*, 65(3):361–393.

Kokoszka, P. (2012). Dependent functional data. *ISRN Probability and Statistics*, 2012.

Kokoszka, P. and Reimherr, M. (2013). Determining the order of the functional autoregressive model. *Journal of Time Series Analysis*, 34(1):116–129.

Korobilis, D. (2013). VAR forecasting using Bayesian variable selection. *Journal of Applied Econometrics*, 28(2):204–230.

Kowal, D. R., Matteson, D. S., and Ruppert, D. (2016). A Bayesian multivariate functional dynamic linear model. *Journal of the American Statistical Association*. (in press).

Kuo, L. and Mallick, B. (1998). Variable selection for regression models. *Sankhyā: The Indian Journal of Statistics, Series B*, pages 65–81.

Laurini, M. P. (2014). Dynamic functional data analysis with non-parametric state space models. *Journal of Applied Statistics*, 41(1):142–163.

Matérn, B. (2013). *Spatial variation*, volume 36. Springer Science & Business Media.

Neal, R. M. (1999). Regression and classification using Gaussian process priors. *Bayesian Statistics*, 6:475–501.

Neal, R. M. (2003). Slice sampling. *Annals of Statistics*, pages 705–741.

Nelson, C. R. and Siegel, A. F. (1987). Parsimonious modeling of yield curves. *Journal of Business*, 60(4):473.

Plummer, M., Best, N., Cowles, K., and Vines, K. (2006). CODA: Convergence diagnosis and output analysis for MCMC. *R News*, 6(1):7–11.

Ramsay, J. O. (2006). *Functional data analysis*. Wiley Online Library.

Ramsay, J. O., Wickham, H., Graves, S., and Hooker, G. (2014). *fda: Functional Data Analysis*. R package version 2.4.4.

Rasmussen, C. E. and Williams, C. K. (2006). Gaussian processes for machine learning. *The MIT Press*.

Shi, J. Q. and Choi, T. (2011). *Gaussian process regression analysis for functional data*. CRC Press.

Wand, M. and Ormerod, J. (2008). On semiparametric regression with O'Sullivan penalized splines. *Australian & New Zealand Journal of Statistics*, 50(2):179–198.

West, M. and Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models*. Springer.

Wood, S. N. (2003). Thin plate regression splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(1):95–114.

Yao, F., Müller, H.-G., and Wang, J.-L. (2005). Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association*, 100(470):577–590.

Appendix

Priors

The prior for $\{\boldsymbol{\mu}_t\}_{t=1}^T$ is determined by (6). Let $\boldsymbol{b}_\psi$ be a J_ψ-dimensional vector of cubic B-spline basis functions with $\min\{|\mathcal{T}_o|/2, 35\} = (J_\psi - 4)$ equally-spaced interior knots. The

tensor product expansion $\psi_\ell(\tau,u) = \boldsymbol{b}'_\psi(\tau)\boldsymbol{\Theta}_{\psi_\ell}\boldsymbol{b}_\psi(u) = \left(\boldsymbol{b}'_\psi(u)\otimes\boldsymbol{b}'_\psi(\tau)\right)\boldsymbol{\theta}_{\psi_\ell}$, where $\boldsymbol{\Theta}_{\psi_\ell}$ is a $J_\psi\times J_\psi$ matrix of unknown coefficients and $\boldsymbol{\theta}_{\psi_\ell} = \text{vec}\left(\boldsymbol{\Theta}_{\psi_\ell}\right)$, is computationally convenient for the FAR surfaces $\{\psi_\ell\}_{\ell=1}^p$. The Gaussian prior $[\boldsymbol{\theta}_{\psi_\ell}|\lambda_{\psi_\ell}] \sim N\left(\mathbf{0}, \lambda_{\psi_\ell}^{-1}\boldsymbol{\Omega}_{\psi_\ell}^{-1}\right)$ induces a Gaussian process prior on ψ_ℓ, where $\boldsymbol{\Omega}_{\psi_\ell}$ is a penalty matrix and λ_{ψ_ℓ} is a smoothing parameter. The standard roughness penalty $\int\int\left\{\frac{\partial^2}{\partial u_1}\psi_\ell(u_1,u_2) + 2\frac{\partial^2}{\partial u_1\partial u_2}\psi_\ell(u_1,u_2) + \frac{\partial^2}{\partial u_2}\psi_\ell(u_1,u_2)\right\}du_1\,du_2$ can be expressed as $\boldsymbol{\theta}'_{\psi_\ell}\boldsymbol{\Omega}_2\boldsymbol{\theta}_{\psi_\ell}$ for a known singular matrix $\boldsymbol{\Omega}_2$. To obtain a proper prior, which is necessary for our model averaging procedure, we combine the roughness penalty with a nonstationarity penalty: a sufficient condition for stationarity of Y_t in model (2) is $\sum_{\ell=1}^p\int\int\psi_\ell^2(\tau,u)\,d\tau\,du < 1$, which can be expressed as $\sum_{\ell=1}^p\boldsymbol{\theta}'_{\psi_\ell}\boldsymbol{\Omega}_0\boldsymbol{\theta}_{\psi_\ell} < 1$ where $\boldsymbol{\Omega}_0$ is a known invertible matrix. We use the prior precision matrix $\boldsymbol{\Omega}_{\psi_\ell} = \boldsymbol{\Omega}_2 + \kappa_\ell\boldsymbol{\Omega}_0$, which penalizes roughness of ψ_ℓ and provides shrinkage toward stationarity, where the trade-off is determined by κ_ℓ. Simulations suggest that the posterior distribution is not sensitive to the choice of κ_ℓ; we fix $\kappa_\ell = 1$ for the simulations and assume $\log(\kappa_\ell) \sim N(0,4)$ for the application. For the smoothing parameter λ_{ψ_ℓ}, we use the half-Cauchy prior of Gelman (2006), which provides excellent mixing of the states $\{s_\ell\}$ in the model averaging procedure. The prior may be expressed hierarchically via the auxiliary variables $\tilde{\lambda}_{\psi_\ell} \sim \text{Gamma}\left(\frac{1}{2},\frac{1}{2}\right)$, $\tilde{\xi}_{\psi_\ell} \sim N(0,10^6)$, and $\tilde{\boldsymbol{\theta}}_{\psi_\ell} \sim N\left(\mathbf{0}, \tilde{\lambda}_{\psi_\ell}^{-1}\boldsymbol{\Omega}_{\psi_\ell}^{-1}\right)$, with the identification $\boldsymbol{\theta}_{\psi_\ell} = \tilde{\xi}_{\psi_\ell}\tilde{\boldsymbol{\theta}}_{\psi_\ell}$ and $\lambda_{\psi_\ell} = \tilde{\xi}_{\psi_\ell}^{-2}\tilde{\lambda}_{\psi_\ell}$.

We use the conditionally conjugate inverse-Gamma priors $\sigma_\nu^{-2}, \sigma_\eta^{-2} \sim \text{Gamma}(10^{-3}, 10^{-3})$ for the measurement error precision and the FDLM approximation error precision, respectively. In some cases, we may prefer smoother sample paths of μ_t, but the paths will not be smooth when σ_η^2 is large. If increasing J_ϵ is infeasible or undesirable, fixing σ_η^2 at some small value, such as $\sigma_\eta^2 = 10^{-6}$, often works well, and can be interpreted as a jitter term for computing a valid inverse of $\boldsymbol{K}_\epsilon$ (Neal, 1999). Assuming the FDLM (7) for the innovation covariance $\boldsymbol{K}_\epsilon$, the factors are distributed $\boldsymbol{e}_t \overset{iid}{\sim} N(\mathbf{0}, \boldsymbol{\Sigma}_e)$ with $\boldsymbol{\Sigma}_e = \text{diag}\left(\{\sigma_j^2\}_{j=1}^{J_\epsilon}\right)$, although many generalizations are available (Kowal et al., 2016). To enforce the ordering constraints $\sigma_1^2 > \sigma_2^2 > \cdots > \sigma_{J_\epsilon}^2 > 0$, recall that the joint distribution (of the precisions) may be written $\left[\sigma_1^{-2},\ldots,\sigma_{J_\epsilon}^{-2}\right] = \left[\sigma_{J_\epsilon}^{-2}\right]\prod_{j=1}^{J_\epsilon-1}\left[\sigma_j^{-2}|\sigma_{j+1}^{-2},\ldots,\sigma_{J_\epsilon}^{-2}\right]$. A noninformative

joint prior that respects the constraints is fully specified by $\sigma_{J_\epsilon}^{-2} \sim \text{Gamma}\,(10^{-3}, 10^{-3})$ and $\left[\sigma_j^{-2} | \sigma_{j+1}^{-2}, \ldots, \sigma_{J_\epsilon}^{-2}\right] = \left[\sigma_j^{-2} | \sigma_{j+1}^{-2}\right] \sim \text{Uniform}\left(0, \sigma_{j+1}^{-2}\right)$ for $j = 1, \ldots, J_\epsilon - 1$. The FLCs are $\phi_j(\tau) = \boldsymbol{b}_\phi'(\tau)\boldsymbol{\xi}_j$, where $\boldsymbol{b}_\phi$ is a low-rank thin plate spline basis with knot locations determined by the quantiles of the observation points, $\mathcal{T}_o$, $\boldsymbol{\xi}_j \sim N(\mathbf{0}, \boldsymbol{\Lambda}_{\phi_j})$, and $\boldsymbol{\Lambda}_{\phi_j}^{-1}$ is the low-rank thin plate spline penalty matrix. We follow Wand and Ormerod (2008) in the singular value decomposition-based diagonalization of the penalty matrix, so that $\boldsymbol{\Lambda}_\phi = \text{diag}\left(10^8, 10^8, \lambda_{\phi_j}^{-1}, \ldots, \lambda_{\phi_j}^{-1}\right)$, which places a noninformative prior on the constant and linear components of the thin plate spline basis, which are unpenalized. The prior precision λ_{ϕ_j} is common among the nonlinear components, and corresponds to the smoothing parameter for the regression function ϕ_j. Following Gelman (2006), we place uniform priors on the standard deviations $\lambda_{\phi_j}^{-1/2} \sim \text{Uniform}\,(0, 10^4)$, which implies the prior for the precision $[\lambda_{\phi_j}] \propto \lambda_{\phi_j}^{-3/2}\mathbf{1}\{\lambda_{\phi_j} > 10^{-8}\}$. The upper bound for the prior standard deviation is selected to match the noninformative components of $\boldsymbol{\Lambda}_{\phi_j}$. The orthonormality constraint is enforced during sampling, which we discuss in Section A of the web supplement. We assume the same parametrization and prior distribution for the mean function, $\mu(\tau) = \boldsymbol{b}_\phi'(\tau)\boldsymbol{\theta}_\mu$.

Proof of Theorem 1

To prove Theorem 1, we use the following well-known results:

Proposition 2. *For random vectors $\boldsymbol{\delta}$ and $\boldsymbol{Y}$ with known mean and covariance, the* unique best linear predictor *of $\boldsymbol{\delta}$ given $\boldsymbol{Y}$ is $\mathbb{E}_{\mathbb{G}}[\boldsymbol{\delta}|\boldsymbol{Y}]$, where $\mathbb{E}_{\mathbb{G}}$ is the expectation computed under the assumption that $(\boldsymbol{\delta}', \boldsymbol{Y}')'$ is jointly Gaussian.*

Proposition 3 (West and Harrison, 1997)**.** *Under a DLM such as model* (6)*, the random vectors $\boldsymbol{y}_{1:T} = (\boldsymbol{y}_1', \ldots, \boldsymbol{y}_T')'$ and $\boldsymbol{\mu}_{1:T} = (\boldsymbol{\mu}_1', \ldots, \boldsymbol{\mu}_T')'$ are jointly Gaussian, conditional on the remaining parameters. In addition, all conditionals and marginals of the joint distribution of $(\boldsymbol{y}_{1:T}', \boldsymbol{\mu}_{1:T}')'$ are Gaussian.*

Note that we could extend $\boldsymbol{\mu}_{1:T}$ to include $\boldsymbol{\mu}$, which is also a Gaussian random vector.

Following Propositions 2 and 3, the proof of Theorem 1 is straightforward:

Proof. (Theorem 1) Let $\mathcal{T}_e$ be fixed and finite such that $\mathcal{T}_e \subset \mathcal{T}$. Given this choice of $\mathcal{T}_e$, we can form the DLM (10) with the appropriately modified terms. Similarly, we can form the Gaussian DLM (6). Proposition 3 implies that $(\boldsymbol{y}'_{1:T}, \boldsymbol{\mu}'_{1:T})'$ under model (6) and conditional on $\boldsymbol{\Theta}$ is jointly Gaussian. Therefore, for any $\boldsymbol{\delta}, \boldsymbol{Y} \subseteq \mathcal{D}_T \cup \{\mu_t(\tau) : \tau \in \mathcal{T}_e, t = 1, \ldots, T\}$, i.e., any subvectors of $(\boldsymbol{y}'_{1:T}, \boldsymbol{\mu}'_{1:T})'$, the distribution of $[\boldsymbol{\delta}|\boldsymbol{Y}, \boldsymbol{\Theta}]$ is Gaussian. Proposition (2) implies that $\hat{\boldsymbol{\delta}}(\boldsymbol{Y}|\boldsymbol{\Theta}) \equiv \mathbb{E}[\boldsymbol{\delta}|\boldsymbol{Y}, \boldsymbol{\Theta}]$, computed under the Gaussian DLM (6), is the unique best linear predictor of $[\boldsymbol{\delta}|\boldsymbol{Y}, \boldsymbol{\Theta}]$ under the DLM (10). □

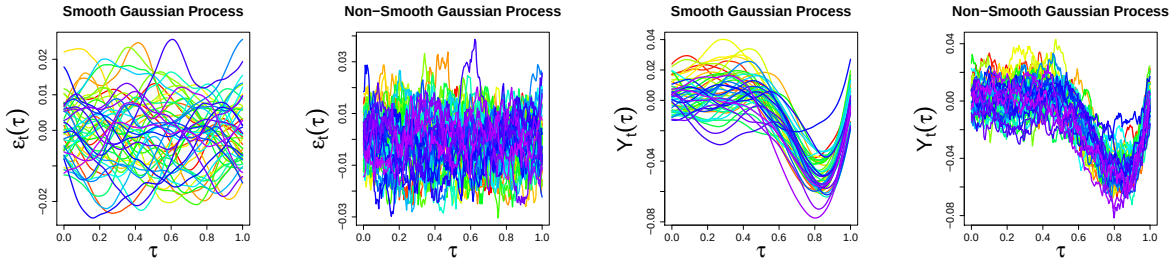

Figure 1: Sample paths of ϵ_t and $Y_t = \mu_t + \mu$ as a function of τ, where ϵ_t is a Gaussian process with the Matérn correlation function, $\boldsymbol{\rho} = (\rho_1, 0.1)$, $\sigma = 0.01$, and Y_t is generated using the Bimodal-Gaussian FAR(1) kernel, $t = 1, \ldots, T = 50$. The curves are time-ordered by color (from red/orange to blue/violet). **Left to right:** $\epsilon_t(\tau), \rho_1 = 2.5$; $\epsilon_t(\tau), \rho_1 = 0.5$; $Y_t(\tau), \rho_1 = 2.5$; $Y_t(\tau), \rho_1 = 0.5$. Note that we do not observe Y_t directly, but rather $y_{i,t} = Y_t(\tau_{i,t}) + \nu_{i,t}$, where $\nu_{i,t} \sim N(0, \sigma_\nu^2)$ is measurement error with $\sigma_\nu = \sigma/5 = 0.002$ and $\mathcal{T}_t = \{\tau_{1,t}, \ldots, \tau_{m_t,t}\}$ are the observation points at time t.

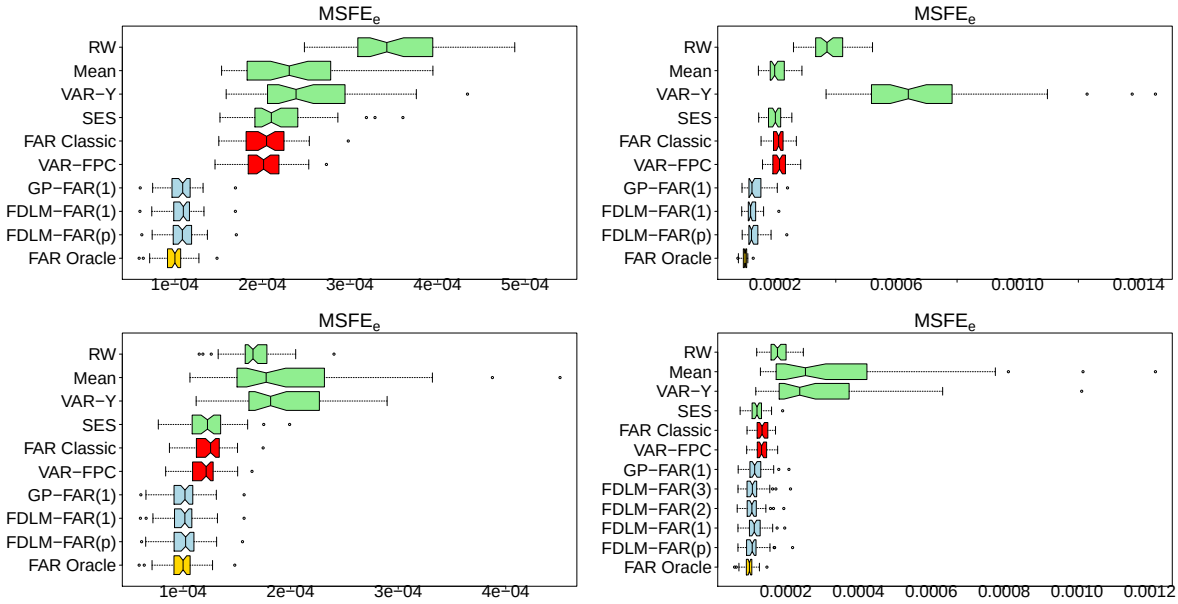

Figure 2: $MSFE_e$ under various designs. **Top left:** FAR(1), $T = 350$, sparse-random design with the Linear-u kernel and smooth GP innovations. **Top right:** FAR(1), $T = 50$, sparse-random design with the Bimodal-Gaussian kernel and non-smooth GP innovations. **Bottom left:** FAR(1), $T = 350$, sparse-fixed design with the Bimodal-Gaussian kernel and smooth GP innovations. **Bottom right:** FAR(2), $T = 125$, sparse-fixed design with Bimodal-Gaussian and Linear$-\tau$ kernels and smooth GP innovations. The proposed methods provide superior forecasts and nearly achieve the oracle performance, despite the presence of sparsity.

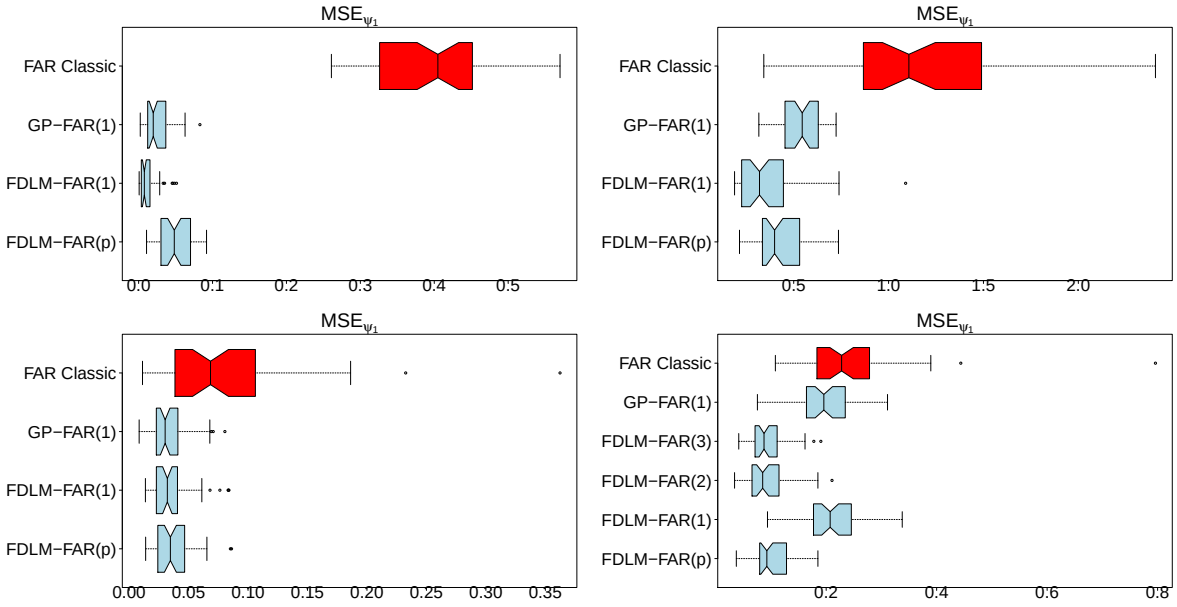

Figure 3: MSE_{ψ_1} under various designs. **Top left:** FAR(1), $T = 350$, sparse-random design with the Linear-u kernel and smooth GP innovations. **Top right:** FAR(1), $T = 50$, sparse-random design with the Bimodal-Gaussian kernel and non-smooth GP innovations. **Bottom left:** FAR(1), $T = 350$, sparse-fixed design with the Bimodal-Gaussian kernel and smooth GP innovations. **Bottom right:** FAR(2), $T = 125$, sparse-fixed design with Bimodal-Gaussian and Linear$-\tau$ kernels and smooth GP innovations. Estimates of ψ_1 are far superior for the proposed methods, including the FAR(p) with model averaging.

Nominal Yields: h-Step Root Mean Squared Forecast Errors (RMSFEs)

	h	RW	Mean	VAR-Y	DL	DRA	FAR Classic	VAR-FPC	FAR(1)	FAR(p)
2/03	1	0.0488	0.4554	*0.0487*	0.1218	0.1440	0.1641	0.1631	0.0498	0.0516
	5	0.0966	0.4369	0.0904	0.1409	0.8221	-	0.1941	*0.0879*	0.1002
8/04	1	0.0253	1.1079	*0.0252*	0.0877	-	0.1113	0.1127	0.0281	0.0281
	5	0.0525	1.1279	*0.0383*	0.0953	-	-	0.1435	0.0412	0.0505
2/06	1	0.1710	0.5408	0.1809	0.2206	-	0.3349	0.3334	0.1682	*0.1673*
	5	*0.4534*	0.5971	0.5885	0.4927	-	-	0.5928	0.4680	0.4627
8/07	1	0.0833	1.3125	0.0860	0.1817	0.1854	0.1168	0.1173	0.0806	*0.0793*
	5	0.1345	1.3146	0.1402	0.2099	0.2998	-	0.1292	0.1537	*0.1233*
2/09	1	*0.0487*	0.5268	0.0517	0.1376	0.0917	0.1406	0.1398	0.0488	0.0760
	5	*0.0894*	0.5560	0.1227	0.1872	0.1451	-	0.1990	0.1323	0.2608
8/10	1	0.0344	0.5063	0.0333	0.1920	0.0878	0.0551	0.0554	*0.0291*	0.0292
	5	0.0583	0.4999	0.0603	0.1950	0.1356	-	0.0724	*0.0452*	0.0495
2/12	1	0.0383	0.5329	0.0384	0.0953	0.1915	0.0464	0.0463	0.0312	*0.0311*
	5	0.0951	0.5522	0.0915	0.1240	0.2476	-	0.0989	0.0760	*0.0734*
8/13	1	0.0463	0.4169	*0.0443*	0.0621	0.0692	0.0634	0.0644	0.0547	0.0676
	5	0.1210	0.3842	0.1104	0.1423	0.1448	-	*0.1100*	0.1208	*0.1100*
2/15	1	0.0329	0.3085	0.0320	0.1125	0.1001	0.0594	0.0606	*0.0305*	0.0321
	5	0.0420	0.3080	0.0403	0.1149	0.1202	-	0.0697	*0.0393*	0.0441

Table 1: h-step RMSFEs for nominal yields, grouped (left to right) by multivariate methods, parametric yield curve models, existing functional data methods, and proposed hierarchical FAR methods. The minimum RMSFE in each row is italicized.

Real Yields: h-Step Root Mean Squared Forecast Errors (RMSFEs)

	h	RW	Mean	VAR-Y	DL	DRA	FAR Classic	VAR-FPC	FAR(1)	FAR(p)
2/03	1	*0.0490*	0.1629	0.0504	0.0499	0.0492	0.1366	0.1329	0.0509	0.0572
	5	0.1001	0.1585	0.1040	0.1017	0.1128	-	0.1525	*0.0967*	0.1110
8/04	1	0.0331	0.3827	0.0337	0.0353	0.0528	0.0431	0.0440	0.0331	*0.0326*
	5	0.0724	0.3924	0.0707	0.0792	0.1690	-	0.0721	0.0679	*0.0651*
2/06	1	0.0429	0.1089	0.0428	0.0448	0.0453	0.0529	0.0533	*0.0424*	*0.0424*
	5	0.0934	0.1082	0.0858	0.0957	0.1362	-	0.0920	0.0852	*0.0835*
8/07	1	*0.0802*	0.2150	0.0896	0.0944	0.1979	0.1212	0.1202	0.0898	0.0880
	5	*0.1866*	0.2309	0.2268	0.2504	1.1843	-	0.1916	0.2051	0.1980
2/09	1	*0.0519*	0.5162	0.0544	0.0643	0.1229	0.0736	0.0749	0.0526	0.0541
	5	*0.0798*	0.5262	0.1100	0.1092	0.3606	-	0.1092	0.0992	0.1046
8/10	1	0.0490	0.7836	0.0492	0.0591	0.0663	0.0800	0.0762	0.0488	*0.0486*
	5	0.0735	0.7845	0.0787	0.0794	0.1815	-	0.0959	*0.0727*	0.0744
2/12	1	*0.0602*	0.8838	0.0612	0.0675	0.1492	0.0906	0.0853	0.0610	0.0608
	5	0.1845	0.9250	0.1958	0.1897	1.7442	-	0.2034	*0.1840*	0.1846
8/13	1	0.0526	0.3242	0.0506	0.0736	-	0.0613	0.0610	0.0500	*0.0492*
	5	0.1551	0.2981	0.1278	0.1380	-	-	0.1246	0.1407	*0.1239*
2/15	1	0.0328	0.3088	0.0327	0.0439	0.1529	0.0776	0.0779	*0.0325*	0.0336
	5	0.0489	0.3104	0.0521	0.0562	-	-	0.0816	*0.0466*	0.0543

Table 2: h-step RMSFEs for real yields, grouped (left to right) by multivariate methods, parametric yield curve models, existing functional data methods, and proposed hierarchical FAR methods. The minimum RMSFE in each row is italicized.

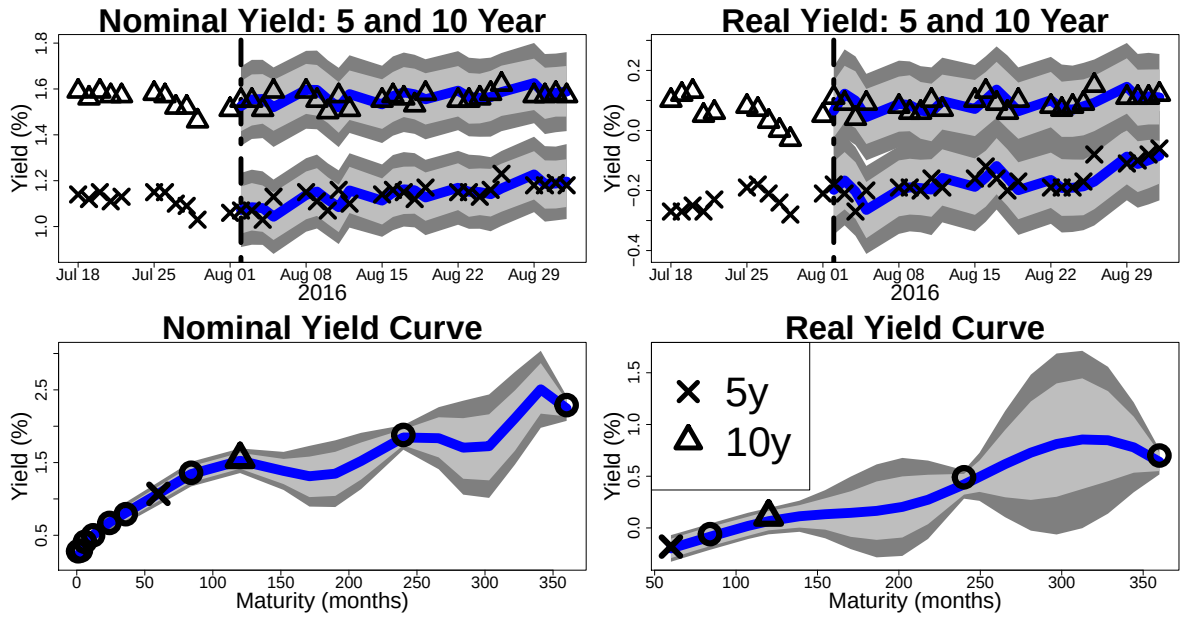

Figure 4: One-step nominal (**left**) and real (**right**) yield curve forecasts during 2016. **Top:** Time series of five (×) and ten (△) year observed maturities with one-step forecasts. **Bottom:** Observed (points) and forecast (line) curves on 8/2/16, corresponding to the dotted vertical line in the top panels. Posterior means (blue) and 95% pointwise and simultaneous prediction bands (light gray and dark gray, respectively) estimated using 10,000 MCMC simulations after a burn-in of 5,000.

Web Supplement

A Initialization and MCMC Sampling Algorithm

A.1 Initialization

We initialize the unknown functions using splines and the remaining parameters using conditional maximum likelihood estimators. We first estimate $\boldsymbol{\mu}$ as a smooth mean of $\{\boldsymbol{y}_t\}_{t=1}^T$, evaluated at $\mathcal{T}_e$. Next, we estimate each $\boldsymbol{\mu}_t$ by fitting a spline to $\boldsymbol{y}_t - \boldsymbol{Z}_t\boldsymbol{\mu}$ for $t = 1, \ldots, T$ using the `R` function `smooth.spline`. Since sparse observation points may lead to unstable initializations of $\boldsymbol{\mu}_t$, we compute the median degrees of freedom implied by the spline fits for $t = 1, \ldots, T$, and then recompute the splines for $t = 1, \ldots, T$ using this common degrees of freedom parameter. Conditional on these estimates, we estimate σ_ν^2, $\{\boldsymbol{\theta}_{\psi_1}, \ldots, \boldsymbol{\theta}_{\psi_p}\}$, and $\{\lambda_{\psi_1}, \ldots, \lambda_{\psi_p}\}$ using the maximum likelihood estimators, and initialize $\tilde{\boldsymbol{\theta}}_{\psi_\ell} = \boldsymbol{\theta}_{\psi_\ell}$, $\tilde{\lambda}_{\psi_\ell} = 1$, and $\tilde{\xi}_{\psi_\ell} = \lambda_{\psi_\ell}^{-1/2}$. From these estimators, we compute the innovations $\boldsymbol{\epsilon}_t$ for $t = 1, \ldots, T$. We initialize the FDLM parameters using the initialization algorithm of Kowal et al. (2016) based on the singular value decomposition (SVD) of $(\boldsymbol{\epsilon}_1, \ldots, \boldsymbol{\epsilon}_T)' = \boldsymbol{U}_e\boldsymbol{D}_e\boldsymbol{V}_e'$. For the FLCs, we let $\boldsymbol{\Phi}$ equal the first J_ϵ columns of $\boldsymbol{V}_e$ and then estimate $\boldsymbol{\Xi}$ to minimize $||\boldsymbol{\Phi} - \boldsymbol{B}_\phi\boldsymbol{\Xi}||^2$. For the factors, we let $(\boldsymbol{e}_1, \ldots, \boldsymbol{e}_T)'$ be the first J_ϵ columns of $(\boldsymbol{U}_e\boldsymbol{D}_e)$, and then estimate $\{\sigma_j^2\}$ and σ_η^2 using the conditional maximum likelihood estimators. Since $\sum_{k=1}^j \sigma_k^2 / \sum_k \sigma_k^2$ estimates the proportion of variance of $\boldsymbol{\epsilon}_t$ explained by the first j factors, we set J_ϵ to be the smallest number of factors that explain at least 95% of the variance of $\boldsymbol{\epsilon}_t$. While more sophisticated procedures are available for selecting J_ϵ, such as DIC and marginal likelihood, we find that this simple approach performs well in simulations.

A.2 Gibbs Sampling Algorithm

We propose to sample from the joint posterior distribution using a Gibbs sampler with the following steps:

1. FAR process, Y_t:

   (a) Centered FAR process, μ_t: form the DLM (6) and sample $\left[\{\boldsymbol{\mu}_t\}_{t=1}^T|\cdots\right]$ *jointly* using the state space sample of Durbin and Koopman (2002) implemented in the `R` package `KFAS`.

   (b) Mean function, $\mu(\tau) = \boldsymbol{b}_\phi'(\tau)\boldsymbol{\theta}_\mu$: sample $[\boldsymbol{\theta}_\mu|\cdots] \sim N(\boldsymbol{A}_\mu\boldsymbol{a}_\mu, \boldsymbol{A}_\mu)$ where

$$\boldsymbol{A}_\mu^{-1} = \boldsymbol{\Lambda}_\mu^{-1} + \sigma_\nu^{-2}\sum_{t=1}^T \boldsymbol{B}_\phi'\boldsymbol{Z}_t'\boldsymbol{Z}_t\boldsymbol{B}_\phi,$$
$$\boldsymbol{a}_\mu = \sigma_\nu^{-2}\sum_{t=1}^T \boldsymbol{B}_\phi'\boldsymbol{Z}_t'(\boldsymbol{y}_t - \boldsymbol{Z}_t\boldsymbol{\mu}_t),$$

   and $\boldsymbol{\Lambda}_\mu = \text{diag}\left(10^8, 10^8, \lambda_\mu^{-1}, \ldots, \lambda_\mu^{-1}\right)$. We sample the smoothing parameter $[\lambda_\mu|\cdots] \sim \text{Gamma}\left(\frac{1}{2}(J_\mu - 3), \frac{1}{2}\sum_{j=3}^{J_\mu}\theta_{\mu,j}^2\right)$ restricted to $\lambda_\mu > 10^{-8}$ (see the σ_j^{-2} sampler below), where J_μ $(= J_\phi)$ is the dimension of $\boldsymbol{\theta}_\mu$ and $\theta_{\mu,j}$ is the jth component of $\boldsymbol{\theta}_\mu$.

   Set $Y_t = \mu_t + \mu$ or, in vector form, $\boldsymbol{Y}_t = \boldsymbol{\mu}_t + \boldsymbol{\mu}$.

2. Measurement error precision, σ_ν^{-2}: sample

$$[\sigma_\nu^{-2}|\cdots] \sim \text{Gamma}\left(10^{-3} + \frac{1}{2}\sum_{t=1}^T m_t, 10^{-3} + \frac{1}{2}\sum_{t=1}^T\sum_{i=1}^{m_t}(y_{i,t} - \mu(\tau_{i,t}) - \mu_t(\tau_{i,t}))^2\right).$$

3. The FAR kernels, $\psi_1, \ldots, \psi_p$: using the Gelman (2006) prior and parametrization of $\boldsymbol{\theta}_{\psi_\ell} = \tilde{\xi}_{\psi_\ell}\tilde{\boldsymbol{\theta}}_{\psi_\ell}$, where $\psi_\ell(\tau, u) = \boldsymbol{b}_\psi'(\tau, u)\boldsymbol{\theta}_{\psi_\ell}$ and $\boldsymbol{B}_\psi = (\boldsymbol{b}_\psi(\tau_1), \ldots, \boldsymbol{b}_\psi(\tau_M))'$, we sample

(a) $\tilde{\boldsymbol{\theta}}_\psi = (\tilde{\boldsymbol{\theta}}'_{\psi_1}, \ldots, \tilde{\boldsymbol{\theta}}'_{\psi_p})'$ *jointly* from $[\tilde{\boldsymbol{\theta}}_\psi | \cdots] \sim N(\boldsymbol{A}_\psi \boldsymbol{a}_\psi, \boldsymbol{A}_\psi)$, where

$$\boldsymbol{A}_\psi^{-1}[\ell, k] = \begin{cases} \lambda_{\psi_\ell} \boldsymbol{\Omega}_{\psi_\ell} + s_\ell \tilde{\xi}_{\psi_\ell}^2 \left[(\boldsymbol{B}'_\psi \boldsymbol{Q}) \left\{ \sum_{t=p+1}^T \boldsymbol{\mu}_{t-\ell} \boldsymbol{\mu}'_{t-\ell} \right\} (\boldsymbol{B}'_\psi \boldsymbol{Q})' \right] \otimes \left[\boldsymbol{B}'_\psi \boldsymbol{K}_\epsilon^{-1} \boldsymbol{B}_\psi \right], & \ell = k \\ s_\ell s_k \tilde{\xi}_{\psi_\ell} \tilde{\xi}_{\psi_k} \left[(\boldsymbol{B}'_\psi \boldsymbol{Q}) \left\{ \sum_{t=p+1}^T \boldsymbol{\mu}_{t-\ell} \boldsymbol{\mu}'_{t-k} \right\} (\boldsymbol{B}'_\psi \boldsymbol{Q})' \right] \otimes \left[\boldsymbol{B}'_\psi \boldsymbol{K}_\epsilon^{-1} \boldsymbol{B}_\psi \right] & \ell \neq k \end{cases}$$

$$\boldsymbol{a}_\psi[\ell] = s_\ell \tilde{\xi}_{\psi_\ell} \text{vec} \left(\boldsymbol{B}'_\psi \boldsymbol{K}_\epsilon^{-1} \left\{ \sum_{t=p+1}^T \boldsymbol{\mu}_t \boldsymbol{\mu}'_{t-\ell} \right\} (\boldsymbol{B}'_\psi \boldsymbol{Q})' \right),$$

$\boldsymbol{A}_\psi^{-1}[\ell, k]$ is the (ℓ, k)th block of $\boldsymbol{A}_\psi^{-1}$ of dimension $J_\psi^2 \times J_\psi^2$ and $\boldsymbol{a}_\psi[\ell]$ is the ℓth subvector of $\boldsymbol{a}_\psi$ of length J_ψ^2;

(b) For $\ell = 1, \ldots, p$, sample $\left[\tilde{\xi}_{\psi_\ell} | \cdots \right] \sim N\left(A_{\tilde{\xi}_{\psi_\ell}} a_{\tilde{\xi}_{\psi_\ell}}, A_{\tilde{\xi}_{\psi_\ell}} \right)$, where

$$A_{\tilde{\xi}_{\psi_\ell}}^{-1} = 10^{-6} + \tilde{\boldsymbol{\theta}}'_\psi \left(\left[(\boldsymbol{B}'_\psi \boldsymbol{Q}) \left\{ \sum_{t=p+1}^T \boldsymbol{\mu}_{t-\ell} \boldsymbol{\mu}'_{t-\ell} \right\} (\boldsymbol{B}'_\psi \boldsymbol{Q})' \right] \otimes \left[\boldsymbol{B}'_\psi \boldsymbol{K}_\epsilon^{-1} \boldsymbol{B}_\psi \right] \right) \tilde{\boldsymbol{\theta}}_\psi,$$

$$a_{\tilde{\xi}_{\psi_\ell}} = \tilde{\boldsymbol{\theta}}'_\psi \text{vec} \left(\boldsymbol{B}'_\psi \boldsymbol{K}_\epsilon^{-1} \left\{ \sum_{t=p+1}^T \left[\boldsymbol{\mu}_t - \sum_{k \neq \ell} s_k \boldsymbol{G}(\psi_k) \boldsymbol{\mu}_{t-k} \right] \boldsymbol{\mu}'_{t-\ell} \right\} (\boldsymbol{B}'_\psi \boldsymbol{Q})' \right),$$

sample $\left[\tilde{\lambda}_{\psi_\ell} | \cdots \right] \sim \text{Gamma}\left(\frac{1}{2} + J_\psi^2/2, \frac{1}{2} + \boldsymbol{\theta}'_{\psi_\ell} \boldsymbol{\Omega}_{\psi_\ell} \boldsymbol{\theta}_{\psi_\ell}/2 \right)$, and, if κ_ℓ is unknown, sample κ_ℓ using the slice sampler (Neal, 2003). Set $\boldsymbol{\theta}_{\psi_\ell} = \tilde{\xi}_{\psi_\ell} \tilde{\boldsymbol{\theta}}_{\psi_\ell}$ and update $\boldsymbol{\Omega}_{\psi_\ell}$.

(c) For the model averaging procedure, sample $[s_\ell | \cdots]$ (in random order), i.e., set $s_\ell = 1$ if $\log O_{10}^{post} > \log(1/U - 1)$ and $s_\ell = 0$ otherwise, where $U \sim \text{Uniform}(0, 1)$, $\log O_{10}^{post}$ is the log-posterior odds

$$\log O_{10}^{post} = -\frac{1}{2} \sum_{t=p+1}^T \left[\boldsymbol{\mu}'_{t-\ell} \boldsymbol{K}_\epsilon^{-1} \boldsymbol{\mu}_{t-\ell} - 2 \left(\boldsymbol{\mu}_t - \sum_{k \neq \ell} s_k \boldsymbol{G}(\psi_k) \boldsymbol{\mu}_{t-k} \right)' \boldsymbol{K}_\epsilon^{-1} \boldsymbol{\mu}_{t-\ell} \right] + \log O_{10}^{prior},$$

and $\log O_{10}^{prior} = \log \mathbb{P}(s_\ell = 1 | s_k, k \neq \ell) - \log \mathbb{P}(s_\ell = 0 | s_k, k \neq \ell)$ is the log-prior odds.

4. The innovation covariance, $\boldsymbol{K}_\epsilon$, under the FDLM:

(a) The factors, $\{\boldsymbol{e}_t\}_{t=1}^T$: using the prior $\boldsymbol{e}_t \overset{iid}{\sim} N(\boldsymbol{0}, \boldsymbol{\Sigma}_e)$ and the conditional likelihood

$\boldsymbol{\epsilon}_t = \boldsymbol{\mu}_t - \sum_{\ell=1}^{p} \boldsymbol{G}(\psi_\ell)\boldsymbol{\mu}_{t-\ell} = \boldsymbol{\Phi}\boldsymbol{e}_t + \boldsymbol{\eta}_t$, sample $[\boldsymbol{e}_t|\cdots] \sim N(\boldsymbol{A}_e\boldsymbol{a}_{e_t}, \boldsymbol{A}_e)$, where

$$\boldsymbol{A}_e^{-1} = \sigma_\eta^{-2}\boldsymbol{\Phi}'\boldsymbol{\Phi} + \boldsymbol{\Sigma}_e^{-1} = \text{diag}\left(\{\sigma_\eta^{-2} + \sigma_j^{-2}\}_{j=1}^{J_\epsilon}\right)$$
$$\boldsymbol{a}_{e_t} = \sigma_\eta^{-2}\boldsymbol{\Phi}'\boldsymbol{\epsilon}_t.$$

Note that $\boldsymbol{A}_e$ is time-invariant and diagonal, so we can sample $\{\boldsymbol{e}_t\}_{t=1}^T$ jointly and efficiently.

(b) The factor precisions, σ_j^{-2}: sample $[\sigma_{J_\epsilon}^{-2}|\cdots] \sim \text{Gamma}\left(10^{-3} + \frac{T}{2}, 10^{-3} + \frac{1}{2}\sum_{t=1}^{T} e_{J_\epsilon,t}^2\right)$; then, for $j = J_\epsilon - 1, \ldots, 1$, set $\sigma_j^{-2} = F_\phi^{-1}(U; s_\phi, r_{\phi_j})$, where F_ϕ is the distribution function for a Gamma random variable with shape parameter $s_\phi = (T-1)/2$ and rate parameter $r_{\phi_j} = \sum_{t=1}^{T} e_{j,t}^2/2$, and $U \sim \text{Uniform}\left(a_{\phi_j}, b_{\phi_j}\right)$ where $a_{\phi_j} = F_\phi(0; s_\phi, r_{\phi_j})$ and $b_{\phi_j} = F_\phi(\sigma_{j+1}^{-2}; s_\phi, r_{\phi_j})$.

(c) The approximation error precision, σ_η^{-2}: sample

$$[\sigma_\eta^{-2}|\cdots] \sim \text{Gamma}\left(10^{-3} + \frac{TM}{2}, 10^{-3} + \frac{1}{2}\sum_{t=1}^{T} ||\boldsymbol{\epsilon}_t - \boldsymbol{\Phi}\boldsymbol{e}_t||^2\right)$$

where $||\cdot||^2$ denotes the Euclidean distance.

(d) The factor loading curves: for $j = 1, \ldots, J_\epsilon$ (in random order), sample $\boldsymbol{\xi}_j \sim N(\boldsymbol{A}_{\xi_j}\boldsymbol{a}_{\xi_j}, \boldsymbol{A}_{\xi_j})$, where

$$\boldsymbol{A}_{\xi_j}^{-1} = \boldsymbol{\Lambda}_{\phi_j}^{-1} + \sigma_\eta^{-2}\left(\sum_{t=1}^{T} e_{j,t}^2\right)\boldsymbol{B}_\phi'\boldsymbol{B}_\phi,$$
$$\boldsymbol{a}_{\xi_j} = \sigma_\eta^{-2}\boldsymbol{B}_\phi'\sum_{t=1}^{T} e_{j,t}\left(\boldsymbol{\epsilon}_t - \boldsymbol{B}_\phi\sum_{k\neq j}\boldsymbol{\xi}_k e_{k,t}\right).$$

To enforce the orthogonality constraint, we condition on the linear constraints $(\boldsymbol{B}_\phi\boldsymbol{\xi}_k)'\boldsymbol{B}_\phi\boldsymbol{\xi}_j = 0$ for $k \neq j$; since $\boldsymbol{\xi}_j$ is Gaussian and $\boldsymbol{\xi}_k$ is conditioned upon, the resulting distribution is Gaussian with easily computable moments, which is also convenient for efficient sampling; see Kowal et al. (2016) for more details. After sam-

pling from the conditional distribution, we normalize the sampled vector $\boldsymbol{\xi}_j$, so that the orthonormality constraint is enforced at every MCMC iteration. We sample the corresponding smoothing parameters $[\lambda_{\phi_j}|\cdots] \sim \text{Gamma}\left(\frac{1}{2}(J_\phi - 3), \frac{1}{2}\sum_{k=3}^{J_\phi}\xi_{j,k}^2\right)$ restricted to $\lambda_{\phi_j} > 10^{-8}$, where $\xi_{j,k}$ is the kth component of $\boldsymbol{\xi}_j$.

Finally, we form the covariance and precision matrices $\boldsymbol{K}_\epsilon$ and $\boldsymbol{K}_\epsilon^{-1}$, respectively, using the sampled components. Since the orthonormality constraint $\boldsymbol{\Phi}'\boldsymbol{\Phi} = \boldsymbol{I}_{J_\epsilon}$ is enforced at every MCMC iteration, we can compute $\boldsymbol{K}_\epsilon^{-1}$ directly and efficiently using (8).

When the sample size T or the number of evaluation points M is large (i.e., $T > 10,000$ or $M > 50$), the Durbin and Koopman (2002) joint sampler is computationally inefficient. Instead, we may use a single-move sampler for $\{\boldsymbol{\mu}_t\}_{t=1}^T$, in which we sample from the full conditional distribution of each $[\boldsymbol{\mu}_t|\boldsymbol{\mu}_s, s \neq t]$ separately for $t = 1, \ldots, T$ (in random order). The single-move sampler is more computationally efficient, but is typically less MCMC efficient. The FDLM provides a closed form for $\boldsymbol{K}_\epsilon^{-1}$, which substantially reduces computation time when M is large.

The tensor product basis for ψ_ℓ provides a computational simplification for jointly sampling the FAR kernel basis coefficients, $\boldsymbol{\theta}_\psi$. Importantly, the dimension of the Kronecker product for computing $\boldsymbol{A}_\psi^{-1}$ is determined by the number of basis functions, J_ψ, which is bounded by 35 in our specification, and may be smaller for some applications. For other bivariate bases, such as the thin plate spline basis, such simplifications are not readily available, and the Kronecker product scales with the number of evaluation points, M.

In the model averaging procedure, there is a nontrivial concern about the ability of the MCMC sampler to move between states. When $s_\ell = 0$, ψ_ℓ does not appear in the likelihood (9), so the Gibbs sampler will draw ψ_ℓ from its prior. Therefore, the prior for ψ_ℓ must be proper; if it is nonetheless noninformative, then the draws of ψ_ℓ from the prior distribution may not be reasonable for (9), so the next MCMC sample of s_ℓ will be zero with high probability. To alleviate this problem, we fix $s_\ell = 1$ for all ℓ during a short burn-in period,

so that each ψ_ℓ is well-estimated and therefore more likely to be included in the model if it is relevant. In both simulations and the yield curve application, the Gelman (2006) parametrization for ψ_ℓ sampling discussed in the Appendix provides excellent mixing among the states $\{s_\ell\}_{\ell=1}^{p_{max}}$.

B Additional Theoretical Results

B.1 Proof of Proposition 1

Let $\Psi(B)$ be a polynomial in the backshift operator B of order p, so that $\Psi(B)Y_t = (1-\Psi_1 B-\Psi_2 B^2 - \cdots - \Psi_p B^p)Y_t = Y_t - \sum_{\ell=1}^{p} \Psi_\ell(Y_{t-\ell})$, where $\{\Psi_\ell\}_{\ell=1}^{p}$ are bounded linear operators on $L^2(\mathcal{T})$. Similarly, let $\Theta(B)$ be a polynomial in the backshift operator B of order q, where $\{\Theta\}_{\ell=1}^{q}$ are bounded linear operators on $L^2(\mathcal{T})$. A *functional autoregressive moving average process* of order (p,q), written FARMA(p,q), is defined by $\Psi(B)(Y_t-\mu) = \Theta(B)\epsilon_t$, where $\{\epsilon_t\}$ is a white noise process in $L^2(\mathcal{T})$ and μ is the unconditional mean of Y_t. The FAR(p) model may be written compactly as $\Psi(B)(Y_t-\mu) = \epsilon_t$. By assumption, we observe the process $\{y_t\}$, where $y_t = Y_t + \nu_t$ and $\{\nu_t\}$ is a white noise process in $L^2(\mathcal{T})$ independent of $\{\epsilon_t\}$. Rewriting the observation equation $y_t - \mu = Y_t - \mu + \nu_t$ and applying $\Psi(B)$, we have $\Psi(B)(y_t-\mu) = \Psi(B)(Y_t-\mu)+\Psi(B)\nu_t = \epsilon_t+\Psi(B)\nu_t$. It remains to show that $Z_t \equiv \epsilon_t+\Psi(B)\nu_t$ is a *functional moving average process* of order p, or equivalently, FARMA$(0,p)$. Clearly, $X_t \equiv \Psi(B)\nu_t$ is FARMA$(0,p)$. By Proposition 10.2 in Bosq and Blanke (2008), $C_p^X \neq 0$ and $C_\ell^X = 0$ for $\ell > p$, where C_ℓ^X is the covariance operator of X_t defined by $C_\ell^X(x) \equiv \mathbb{E}\left[\langle X_t, x\rangle X_{t+\ell}\right]$ for $x \in L^2(\mathcal{T})$. Let C_ℓ^Z and C_ℓ^ϵ denote the covariance operators for Z_t and ϵ_t, respectively. Then $C_\ell^Z(x) = \mathbb{E}\left[\langle Z_t, x\rangle Z_{t+\ell}\right] = \mathbb{E}\left[\langle \epsilon_t + X_t, x\rangle \left(\epsilon_{t+\ell} + X_{t+\ell}\right)\right] = C_\ell^\epsilon(x) + C_\ell^X(x)$, using independence of $\{\epsilon_t\}$ and $\{\nu_t\}$. Since ϵ_t is white noise, $C_\ell^\epsilon = 0$ for $\ell > 0$, from which it follows that $C_p^Z \neq 0$ and $C_\ell^Z = 0$ for $\ell > p$. Proposition 10.2 in Bosq and Blanke (2008) implies that Z_t is FARMA$(0,p)$, so we conclude that y_t is FARMA(p,p).

B.2 DLM Recursions and Special Cases of Theorem 1

For completeness, we provide the standard DLM recursion formulas for model (6). Let $\mathcal{D}_t = \{\boldsymbol{y}_t, \boldsymbol{y}_{t-1}, \ldots, \boldsymbol{y}_1\} \cup \mathcal{D}_0$ be the information available at time t, where $\mathcal{D}_0$ represents the information prior to $t = 1$. For our purposes—in particular, for the Gibbs sampling algorithm—we let $\mathcal{D}_0 = \{\mu, \sigma_\nu^2, \psi, K_\epsilon\}$ (denoted by $\boldsymbol{\Theta}$ in Theorem 1). We may compute full conditional posterior distributions from model (6) using standard DLM recursions (e.g., West and Harrison, 1997). For simplicity, let $\boldsymbol{G} = \boldsymbol{G}(\psi)$. Suppose that $[\boldsymbol{\mu}_{t-1}|\mathcal{D}_{t-1}] \sim N(\boldsymbol{m}_{t-1}, \boldsymbol{C}_{t-1})$. The *prior at time t* is $[\boldsymbol{\mu}_t|\mathcal{D}_{t-1}] \sim N(\boldsymbol{a}_t, \boldsymbol{R}_t)$, where $\boldsymbol{a}_t = \boldsymbol{G}\boldsymbol{m}_{t-1}$ and $\boldsymbol{R}_t = \boldsymbol{G}\boldsymbol{C}_{t-1}\boldsymbol{G}' + \boldsymbol{K}_\epsilon$. The *one-step forecast at time t* is $[\boldsymbol{y}_t|\mathcal{D}_{t-1}] \sim N(\boldsymbol{f}_t, \boldsymbol{Q}_t)$, where $\boldsymbol{f}_t = \boldsymbol{Z}_t\boldsymbol{\mu} + \boldsymbol{Z}_t\boldsymbol{a}_t = \boldsymbol{Z}_t(\boldsymbol{\mu} + \boldsymbol{G}\boldsymbol{m}_{t-1})$ and $\boldsymbol{Q}_t = \boldsymbol{Z}_t\boldsymbol{R}_t\boldsymbol{Z}_t' + \sigma_\nu^2\boldsymbol{I}_{m_t}$. The *posterior at time t* is $[\boldsymbol{\mu}_t|\mathcal{D}_t] \sim N(\boldsymbol{m}_t, \boldsymbol{C}_t)$, where $\boldsymbol{m}_t = \boldsymbol{C}_t^{-1}\left(\boldsymbol{R}_t^{-1}\boldsymbol{a}_t + \sigma_\nu^{-2}\boldsymbol{Z}_t'(\boldsymbol{y}_t - \boldsymbol{Z}_t\boldsymbol{\mu})\right)$ and $\boldsymbol{C}_t^{-1} = \boldsymbol{R}_t^{-1} + \sigma_\nu^{-2}\boldsymbol{Z}_t'\boldsymbol{Z}_t$, or, more commonly, $\boldsymbol{m}_t = \boldsymbol{a}_t + \boldsymbol{A}_t\boldsymbol{r}_t$, $\boldsymbol{A}_t = \boldsymbol{R}_t\boldsymbol{Z}_t'\boldsymbol{Q}_t^{-1}$, $\boldsymbol{r}_t = \boldsymbol{y}_t - \boldsymbol{f}_t$, and $\boldsymbol{C}_t = \boldsymbol{R}_t - \boldsymbol{A}_t\boldsymbol{Q}_t\boldsymbol{A}_t'$. The h-step forecast of the functional observations is $\mathbb{E}[\boldsymbol{y}_{t+h}|\mathcal{D}_t] = \mathbb{E}[\boldsymbol{Z}_{t+h}\boldsymbol{\mu} + \boldsymbol{Z}_{t+h}\boldsymbol{\mu}_{t+h} + \boldsymbol{\nu}_{t+h}|\mathcal{D}_t] = \boldsymbol{Z}_{t+h}\boldsymbol{\mu} + \boldsymbol{Z}_{t+h}\mathbb{E}[\boldsymbol{\mu}_{t+h}|\mathcal{D}_t]$, where $\mathbb{E}[\boldsymbol{\mu}_{t+h}|\mathcal{D}_t] = \boldsymbol{G}^h\boldsymbol{m}_t$, which is the h-step forecast of $\boldsymbol{\mu}_t$.

Some special cases of Theorem 1 are proved in West and Harrison (1997):

Corollary B.2.1 (Theorem 4.10, West and Harrison, 1997). *The unique best linear predictor of the filtering random variable* $[\boldsymbol{\mu}_t|\mathcal{D}_t]$ *is* $\boldsymbol{m}_t$

Corollary B.2.2 (Corollary 4.7, West and Harrison, 1997). *The unique best linear predictor of the one-step forecast* $[\boldsymbol{\mu}_t|\mathcal{D}_{t-1}]$ *is* $\boldsymbol{a}_t$. *The unique best linear predictor of the one-step forecast* $[\boldsymbol{y}_t|\mathcal{D}_{t-1}]$ *is* $\boldsymbol{f}_t$.

B.3 Proof of Theorem 2

Suppose $\tau^* \in \mathcal{T}$ such that $\tau^* \notin \mathcal{T}_e$. The full conditional distribution of $\mu_t(\tau^*)$ is

$$\begin{aligned}\left[\mu_t(\tau^*)|\{\boldsymbol{\mu}_r\}_{r=1}^T, \boldsymbol{\Theta}, \mathcal{D}_s\right] &\propto \left[\boldsymbol{y}_1, \ldots, \boldsymbol{y}_s|\mu_t(\tau^*), \{\boldsymbol{\mu}_r\}_{r=1}^T, \boldsymbol{\Theta}\right] \times \left[\mu_t(\tau^*)|\{\boldsymbol{\mu}_r\}_{r=1}^T, \boldsymbol{\Theta}\right] \\ &\propto \left[\mu_t(\tau^*)|\{\boldsymbol{\mu}_r\}_{r=1}^T, \boldsymbol{\Theta}\right],\end{aligned}$$

since the likelihood term is constant with respect to $\mu_t(\tau^*)$: $\mathcal{T}_o \subseteq \mathcal{T}_e$, so $\tau^* \notin \mathcal{T}_e$ implies $\tau^* \notin \mathcal{T}_o$, and therefore $\mu_t(\tau^*)$ does not appear in the likelihood of model (4). For $p = 1$, the conditional Gaussian process prior for μ_t implied by model (4) under the approximation (5) is $[\mu_t|\mu_{t-1}, \psi, K_\epsilon] \sim \mathcal{GP}\left(\boldsymbol{\psi}'(\cdot)\boldsymbol{Q}\boldsymbol{\mu}_{t-1}, K_\epsilon\right)$, where $\boldsymbol{\psi}'(\tau) = (\psi(\tau, \tau_1), \ldots, \psi(\tau, \tau_M))$, $\boldsymbol{Q}$ is a known quadrature weight matrix, and $\boldsymbol{\mu}_{t-1} = (\mu_{t-1}(\tau_1), \ldots, \mu_{t-1}(\tau_M))'$ is the function μ_{t-1} evaluated at each $\tau \in \mathcal{T}_e$. Notably, $\tau^* \notin \mathcal{T}_e$ implies that $\mu_t(\tau^*)$ does not appear in the conditional mean function for μ_{t+1}, so we may further simplify the distribution of $\mu_t(\tau^*)$:

$$\left[\mu_t(\tau^*)|\{\boldsymbol{\mu}_r\}_{r=1}^T, \boldsymbol{\Theta}, \mathcal{D}_s\right] \propto \left[\mu_t(\tau^*)|\boldsymbol{\mu}_t, \boldsymbol{\mu}_{t-1}, \boldsymbol{\Theta}\right].$$

To compute this distribution, we use the definition of a Gaussian process, which implies the following joint distribution of $\mu_t(\tau^*)$ and $\boldsymbol{\mu}_t$, conditional on $\boldsymbol{\mu}_{t-1}$, ψ, and K_ϵ:

$$\begin{pmatrix} \mu_t(\tau^*) \\ \boldsymbol{\mu}_t \end{pmatrix} \sim N\left(\begin{pmatrix} \boldsymbol{\psi}'(\tau^*)\boldsymbol{Q}\boldsymbol{\mu}_{t-1} \\ \boldsymbol{\Psi}\boldsymbol{Q}\boldsymbol{\mu}_{t-1} \end{pmatrix}, \begin{pmatrix} K_\epsilon(\tau^*, \tau^*) & \boldsymbol{K}_\epsilon(\tau^*) \\ \boldsymbol{K}_\epsilon'(\tau^*) & \boldsymbol{K}_\epsilon \end{pmatrix}\right),$$

where $\boldsymbol{\Psi} = \{\psi(\tau_i, \tau_k)\}_{i,k=1}^M$ and $\boldsymbol{K}_\epsilon(\tau^*) = (K_\epsilon(\tau^*, \tau_1), \ldots, K_\epsilon(\tau^*, \tau_M))$. Conditioning on $\boldsymbol{\mu}_t$ induces the desired distribution $[\mu_t(\tau^*)|\boldsymbol{\mu}_t, \boldsymbol{\mu}_{t-1}, \psi, K_\epsilon] \sim N\left(m_t(\tau^*), K_t(\tau^*)\right)$, where $m_t(\tau^*) = \boldsymbol{\psi}'(\tau^*)\boldsymbol{Q}\boldsymbol{\mu}_{t-1} + \boldsymbol{K}_\epsilon(\tau^*)\boldsymbol{K}_\epsilon^{-1}\left(\boldsymbol{\mu}_t - \boldsymbol{\Psi}\boldsymbol{Q}\boldsymbol{\mu}_{t-1}\right)$ and $K_t(\tau^*) = K_\epsilon(\tau^*, \tau^*) - \boldsymbol{K}_\epsilon(\tau^*)\boldsymbol{K}_\epsilon^{-1}\boldsymbol{K}_\epsilon'(\tau^*)$. Under the FDLM, the following useful simplifications are available: $K_\epsilon(\tau^*, \tau^*) = \sigma_\eta^2 + \boldsymbol{\phi}'(\tau^*)\boldsymbol{\Sigma}_e\boldsymbol{\phi}(\tau^*)$, $\boldsymbol{K}_\epsilon(\tau^*) = \boldsymbol{\phi}'(\tau^*)\boldsymbol{\Sigma}_e\boldsymbol{\Phi}'$, and using (8), $\boldsymbol{K}_\epsilon^{-1} = \sigma_\eta^{-2}\boldsymbol{I}_M - \sigma_\eta^{-2}\boldsymbol{\Phi}\tilde{\boldsymbol{\Sigma}}_e\boldsymbol{\Phi}'$, where $\boldsymbol{\phi}'(\tau^*) = (\phi_1(\tau^*), \ldots, \phi_{J_\epsilon}(\tau^*))$, $\boldsymbol{\Sigma}_e = \text{diag}\left(\{\sigma_j^2\}_{j=1}^{J_\epsilon}\right)$, $\boldsymbol{\Phi} = (\boldsymbol{\phi}(\tau_1), \ldots, \boldsymbol{\phi}(\tau_M))'$, and $\tilde{\boldsymbol{\Sigma}}_e =$

diag $\left(\{\sigma_j^2/(\sigma_\eta^2+\sigma_j^2)\}_{j=1}^{J_\epsilon}\right)$. By substitution, we derive

$$\begin{aligned} m_t(\tau^*) &= \boldsymbol{\psi}'(\tau^*)\boldsymbol{Q}\boldsymbol{\mu}_{t-1} + \boldsymbol{K}_\epsilon(\tau^*)\boldsymbol{K}_\epsilon^{-1}\left(\boldsymbol{\mu}_t - \boldsymbol{\Psi}\boldsymbol{Q}\boldsymbol{\mu}_{t-1}\right) \\ &= \boldsymbol{\psi}'(\tau^*)\boldsymbol{Q}\boldsymbol{\mu}_{t-1} + \boldsymbol{\phi}'(\tau^*)\boldsymbol{\Sigma}_e\boldsymbol{\Phi}'\left(\sigma_\eta^{-2}\boldsymbol{I}_M - \sigma_\eta^{-2}\boldsymbol{\Phi}\tilde{\boldsymbol{\Sigma}}_e\boldsymbol{\Phi}'\right)\left(\boldsymbol{\mu}_t - \boldsymbol{\Psi}\boldsymbol{Q}\boldsymbol{\mu}_{t-1}\right) \\ &= \boldsymbol{\psi}'(\tau^*)\boldsymbol{Q}\boldsymbol{\mu}_{t-1} + \boldsymbol{\phi}'(\tau^*)\tilde{\boldsymbol{\Sigma}}_e\boldsymbol{\Phi}'\left(\boldsymbol{\mu}_t - \boldsymbol{\Psi}\boldsymbol{Q}\boldsymbol{\mu}_{t-1}\right), \end{aligned}$$

using the constraint $\boldsymbol{\Phi}'\boldsymbol{\Phi} = \boldsymbol{I}_{J_\epsilon}$ and the simplification $\sigma_\eta^{-2}\boldsymbol{\Sigma}_e - \sigma_\eta^{-2}\boldsymbol{\Sigma}_e\tilde{\boldsymbol{\Sigma}}_e = \tilde{\boldsymbol{\Sigma}}_e$. Similarly,

$$\begin{aligned} K_t(\tau^*) &= K_\epsilon(\tau^*,\tau^*) - \boldsymbol{K}_\epsilon(\tau^*)\boldsymbol{K}_\epsilon^{-1}\boldsymbol{K}_\epsilon'(\tau^*) \\ &= \sigma_\eta^2 + \boldsymbol{\phi}'(\tau^*)\boldsymbol{\Sigma}_e\boldsymbol{\phi}(\tau^*) - \boldsymbol{\phi}'(\tau^*)\boldsymbol{\Sigma}_e\boldsymbol{\Phi}'\left(\sigma_\eta^{-2}\boldsymbol{I}_M - \sigma_\eta^{-2}\boldsymbol{\Phi}\tilde{\boldsymbol{\Sigma}}_e\boldsymbol{\Phi}'\right)\boldsymbol{\Phi}\boldsymbol{\Sigma}_e\boldsymbol{\phi}(\tau^*) \\ &= \sigma_\eta^2 + \sigma_\eta^2\boldsymbol{\phi}'(\tau^*)\tilde{\boldsymbol{\Sigma}}_e\boldsymbol{\phi}(\tau^*), \end{aligned}$$

which is time-invariant. Extensions for $p > 1$ only require modification of the mean function: $m_t(\tau^*) = \sum_{\ell=1}^p \boldsymbol{\psi}_\ell'(\tau^*)\boldsymbol{Q}\boldsymbol{\mu}_{t-\ell} + \boldsymbol{\phi}'(\tau^*)\tilde{\boldsymbol{\Sigma}}_e\boldsymbol{\Phi}'\left(\boldsymbol{\mu}_t - \sum_{\ell=1}^p \boldsymbol{\Psi}_\ell\boldsymbol{Q}\boldsymbol{\mu}_{t-\ell}\right)$, where $\boldsymbol{\psi}_\ell'(\tau) = (\psi_\ell(\tau,\tau_1),\ldots,\psi_\ell(\tau,\tau_M))$ and $\boldsymbol{\Psi}_\ell = \{\psi_\ell(\tau_i,\tau_k)\}_{i,k=1}^M$.

C Additional Simulation Results

In Figure C1, we display the results from FAR(1) simulations under the dense design, while varying both smoothness of ϵ_t and the sample size, T. The functional data methods all nearly achieve the oracle performance, and are superior to the multivariate methods. These results confirm the findings of Didericksen et al. (2012): when T is large and the observation points are dense in $\mathcal{T}$, existing functional data methods can nearly achieve the oracle performance, even when ψ_1 is estimated poorly. The proposed methods, particularly with the FDLM (FDLM-FAR(1) and FDLM-FAR(p)), outperform existing functional data methods for non-smooth GP innovations, and again are far superior for ψ_1 estimation. The uncertainty of p incorporated into the lag selection procedure (FDLM-FAR(p)) does not appear to inhibit forecasting or estimation of ψ_1 substantially.

For further clarity, we plot the *Bimodal-Gaussian* kernel in Figure C2, which is featured prominently in our simulation study.

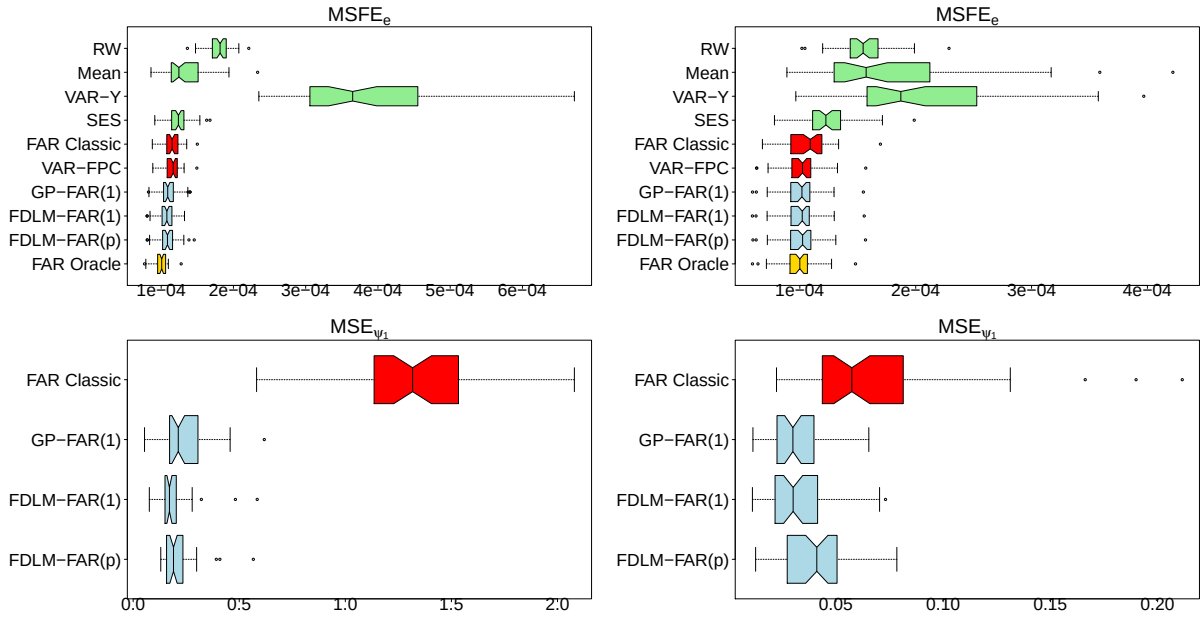

Figure C1: $MSFE_e$ (**top**) and corresponding MSE_{ψ_1} (**bottom**) under various designs. **Left:** FAR(1), $T = 50$, dense design with the Bimodal-Gaussian kernel and non-smooth GP innovations. **Right:** FAR(1), $T = 350$, dense design with the Bimodal-Gaussian kernel and smooth GP innovations. The proposed methods provide superior forecasts and nearly achieve the oracle performance, despite the presence of sparsity.

D Additional Details for the Yield Curve Application

We include MCMC diagnostics for the yield curve application. All diagnostics were computed using the `R` package `coda` (Plummer et al., 2006). In Figures D1 and D2, we provide trace plots for the one-step forecast distributions for the nominal and real yield curves, respectively, on a single day in 2016 across selected maturities. The mixing is very efficient, which is confirmed by effective sample sizes which exceed 5,000 in all cases.

In our yield curve forecasting study of Section 7, we included two popular parametric yield curve models based on the Nelson-Siegel parametrization (Nelson and Siegel, 1987): Diebold and Li (2006, DL) and Diebold et al. (2006, DRA). The Nelson-Siegel basis is defined by

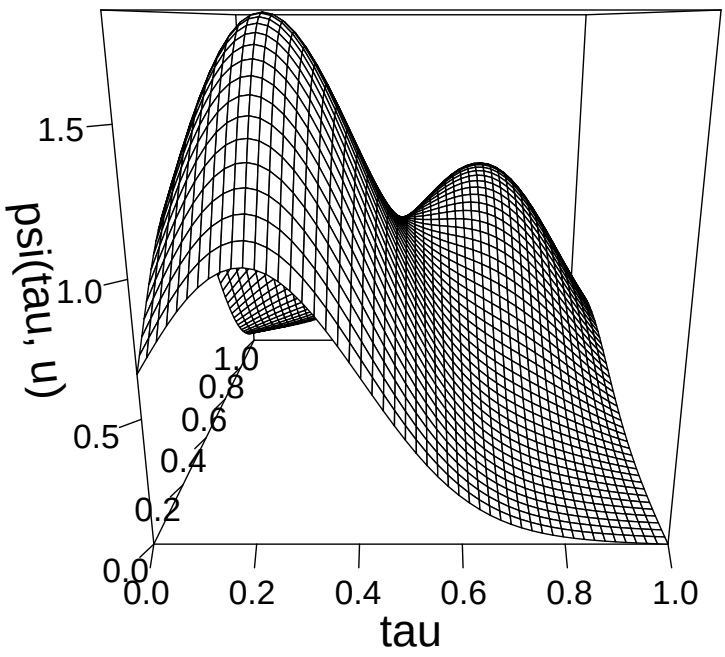

Figure C2: The *Bimodal-Gaussian* kernel, $\psi(\tau,u) \propto \frac{0.75}{\pi(0.3)(0.4)} \exp\{-(\tau-0.2)^2/(0.3)^2 - (u-0.3)^2/(0.4)^2\} + \frac{0.45}{\pi(0.3)(0.4)} \exp\{-(\tau-0.7)^2/(0.3)^2 - (u-0.8)^2/(0.4)^2\}$, normalized so that $\int\int \psi_\ell^2(\tau,u)\, d\tau\, du = 0.8$.

$f_1(\tau) = 1$, $f_2(\tau|\lambda_{NS}) = \frac{1-\exp(-\tau\lambda_{NS})}{\tau\lambda_{NS}}$, and $f_3(\tau|\lambda_{NS}) = \frac{1-\exp(-\tau\lambda_{NS})}{\tau\lambda_{NS}} - \exp(-\tau\lambda_{NS})$, where λ_{NS} is an unknown parameter. For both DL and DRA, the yield curve $Y_t(\tau)$ for time t and time to maturity τ is written as a linear combination of the Nelson-Siegel basis function, for which the corresponding weights are dynamic:

$$Y_t(\tau) = \boldsymbol{f}'(\tau|\lambda_{NS})\boldsymbol{\beta}_t + \epsilon_t(\tau), \tag{11}$$

$$(\boldsymbol{\beta}_t - \boldsymbol{\mu}_\beta) = \boldsymbol{A}\,(\boldsymbol{\beta}_{t-1} - \boldsymbol{\mu}_\beta) + \boldsymbol{\eta}_t \tag{12}$$

where $\boldsymbol{f}'(\tau|\lambda_{NS}) = (f_1(\tau), f_2(\tau|\lambda_{NS}), f_3(\tau|\lambda_{NS}))$, $\boldsymbol{\beta}_t$ is the corresponding 3-dimensional vec-

tor of dynamic weights with unconditional mean $\boldsymbol{\mu}_\beta$, and $\boldsymbol{A}$ is the 3×3 evolution matrix. For implementation purposes, assume that the yield curve is observed at a fixed set of maturities $\tau_1, \ldots, \tau_M$, so that (11) becomes

$$\boldsymbol{y}_t = \boldsymbol{F}_{NS}\boldsymbol{\beta}_t + \boldsymbol{\epsilon}_t \tag{13}$$

where $\boldsymbol{y}_t = (Y_t(\tau_1), \ldots, Y_t(\tau_M))'$, $\boldsymbol{F}_{NS} = (\boldsymbol{f}(\tau_1|\lambda_{NS}), \ldots, \boldsymbol{f}(\tau_M|\lambda_{NS}))'$, and $\boldsymbol{\epsilon}_t = (\epsilon_t(\tau_1), \ldots, \epsilon_t(\tau_M))'$.

The DL approach fixes $\lambda_{NS} = 0.0609$ and then estimates the parameters using a multi-step procedure. First, the weights $\{\boldsymbol{\beta}_t\}$ are estimated using ordinary least squares from (13). Next, the evolution matrix $\boldsymbol{A}$ in (12) is estimated as a VAR coefficient matrix, conditional on $\{\boldsymbol{\beta}_t\}$. Diebold and Li (2006) note that constraining $\boldsymbol{A}$ to be diagonal may improve forecasting in some cases. Finally, h-step forecasts $\hat{\boldsymbol{y}}_{T+h}$ are computed via $\hat{\boldsymbol{y}}_{T+h} = \boldsymbol{F}_{NS}\hat{\boldsymbol{\beta}}_{T+h}$, where $\hat{\boldsymbol{\beta}}_{T+h}$ is the h-step forecast computed from the VAR in (12).

Alternatively, the DRA approach combines (13) and (12) into a state space model, with error distributions $\boldsymbol{\epsilon}_t \overset{iid}{\sim} N(\boldsymbol{0}, \boldsymbol{H})$ independent of $\boldsymbol{\eta}_t \overset{iid}{\sim} N(\boldsymbol{0}, \boldsymbol{Q})$. DRA assume that $\boldsymbol{H}$ is diagonal; we further assume that $\boldsymbol{Q}$ is diagonal, which helps stabilize computations. The unknown parameters $\{\lambda_{NS}, \boldsymbol{A}, \boldsymbol{H}, \boldsymbol{Q}\}$ are then estimated jointly using maximum likelihood based on the Kalman filter. Following DRA, we model λ_{NS} and the diagonal elements of $\boldsymbol{H}$ and $\boldsymbol{Q}$ on the log-scale to ensure positivity in the optimization routine. Conditional on the maximum likelihood estimates for these parameters, DRA use standard state space computations to construct forecasts for the response vector, $\boldsymbol{y}_t$.

E Additional Details on the Quadrature Approximation

Consider the integral in the FAR(1) evolution equation, $\mathcal{I}(\tau) \equiv \int \psi(\tau, u)\mu_{t-1}(u)\, du$, where we omit dependence of $\mathcal{I}$ on t for notational simplicity. In the proposed methodology, we approximate this integral using quadrature: $\mathcal{I}(\tau) \approx \mathcal{I}_M(\tau) \equiv (\psi(\tau, \tau_1), \ldots, \psi(\tau, \tau_M))\, \boldsymbol{Q}\boldsymbol{\mu}_{t-1}$, where $\{\tau_1, \ldots, \tau_M\} = \mathcal{T}_e \subset \mathcal{T}$ is the set of unique evaluation points, $\boldsymbol{Q}$ is a known $M \times M$ quadrature matrix, and $\boldsymbol{\mu}_{t-1} = (\mu_{t-1}(\tau_1), \ldots, \mu_{t-1}(\tau_M))'$ is the function μ_{t-1} evaluated at the

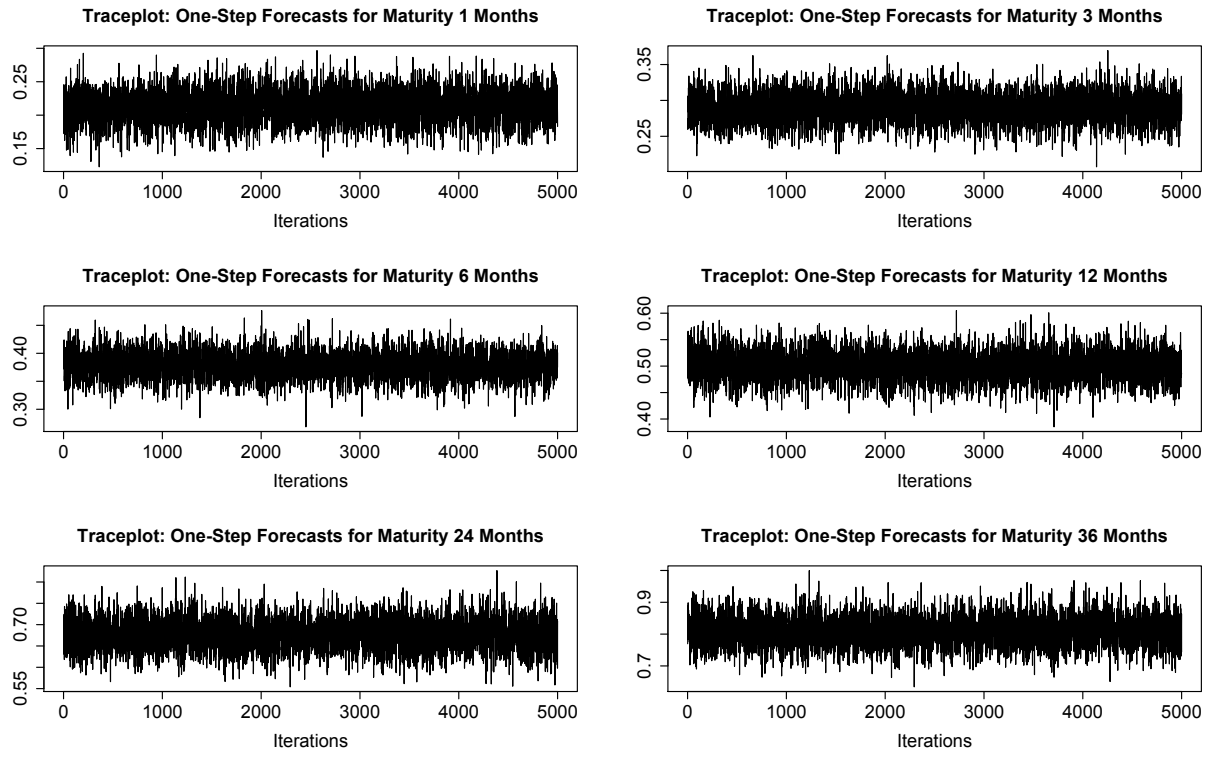

Figure D1: Traceplot for one-step forecasts for nominal yield curves at selected maturities during 2016.

evaluation points. It is important to assess how the accuracy of the approximation of $\mathcal{I}$ by $\mathcal{I}_M$ depends in M, and in particular to determine a value of M sufficiently large to produce reasonable approximations in practice. However, there is a tradeoff: the state vector in the dynamic linear model is M-dimensional, so increasing M indiscriminately may unnecessarily increase computation time.

We conducted a sensitivity analysis based on the simulations from Section 6 of the main paper. In particular, we use the Bimodal-Gaussian kernel, $\psi(\tau, u) \propto \frac{0.75}{\pi(0.3)(0.4)} \exp\{-(\tau - 0.2)^2/(0.3)^2 - (u - 0.3)^2/(0.4)^2\} + \frac{0.45}{\pi(0.3)(0.4)} \exp\{-(\tau - 0.7)^2/(0.3)^2 - (u - 0.8)^2/(0.4)^2\}$, normalized so that $\int\int \psi_\ell^2(\tau, u)\, d\tau\, du = 0.8$. The Bimodal-Gaussian kernel is nonlinear, and therefore is inherently more difficult to approximate using linear quadrature methods, such as the trapezoidal rule. For the other component of the integrand, μ_{t-1}, we simulate $\mu_{t-1} \sim \mathcal{GP}(0, K_\epsilon)$ using the covariance function parameterization $K_\epsilon = \sigma^2 R_\rho$, where R_ρ is

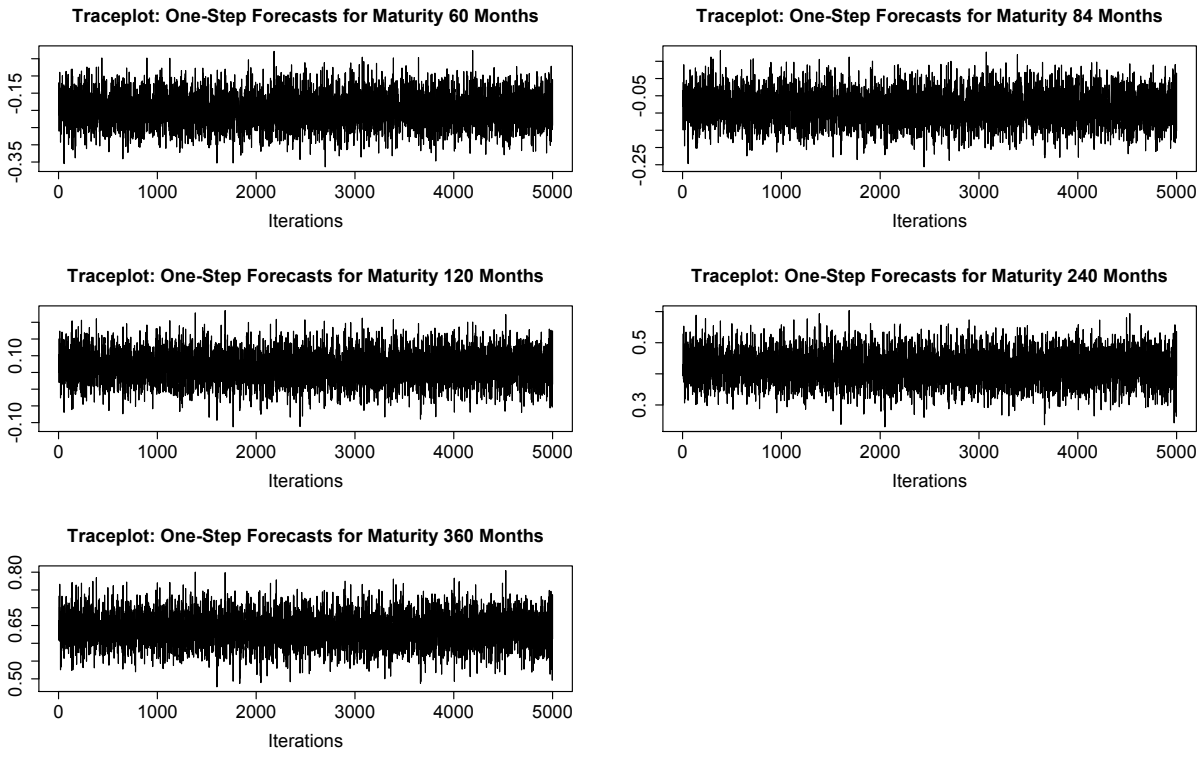

Figure D2: Traceplot for one-step forecasts for real yield curves at selected maturities during 2016.

the Matérn correlation function $R_\rho(\tau,u) = \{2^{\rho_1-1}\Gamma(\rho_1)\}^{-1}\left(||\tau-u||/\rho_2\right)^{\rho_1} K_{\rho_1}(||\tau-u||/\rho_2)$, $\Gamma(\cdot)$ is the gamma function, K_{ρ_1} is the modified Bessel function of order ρ_1, and $\boldsymbol{\rho} = (\rho_1,\rho_2)$ are parameters (Matérn, 2013). We let $\sigma = 0.01$ and $\boldsymbol{\rho} = (\rho_1, 0.1)$, with $\rho_1 = 2.5$ for smooth (twice-differentiable) sample paths and $\rho_1 = 0.5$ for non-smooth (continuous, non-differentiable) sample paths. Comparisons between these cases are important: the non-smooth setting is substantially more challenging for approximations.

For each simulated value of $\mu_{t-1} \sim \mathcal{GP}(0, K_\epsilon)$, we compute $\mathcal{I}_{200}(\tau)$, which we use as a proxy for the true (but unknown) integral value $\mathcal{I}(\tau)$, and compare it to $\mathcal{I}_M(\tau)$ for $M \in \{5,10,15,20,25,30,40,50,60,70,80,90,100\}$. Note that the approximation induced by $\mathcal{I}_{200}(\tau)$ is also used to generate the simulations of Section 6 in the main paper. We measure accuracy using the *relative absolute error* (RAE) and the *standardized squared error*

(SSE), defined respectively by

$$R_M = \int \left| \frac{\mathcal{I}_{200}(\tau) - \mathcal{I}_M(\tau)}{\mathcal{I}_{200}(\tau)} \right| d\tau, \quad S_M = \int \frac{\left(\mathcal{I}_{200}(\tau) - \mathcal{I}_M(\tau)\right)^2}{\sigma^2} d\tau, \tag{14}$$

which we compute for each simulation. We report the pointwise medians for each R_M and S_M as a function of M in Figure E1. As expected, for fixed M, the integral approximation is more accurate when μ_{t-1}—and therefore the integrand—is smooth. Nonetheless, the relative gains of increasing M decline quickly for $M > 20$ in both cases.

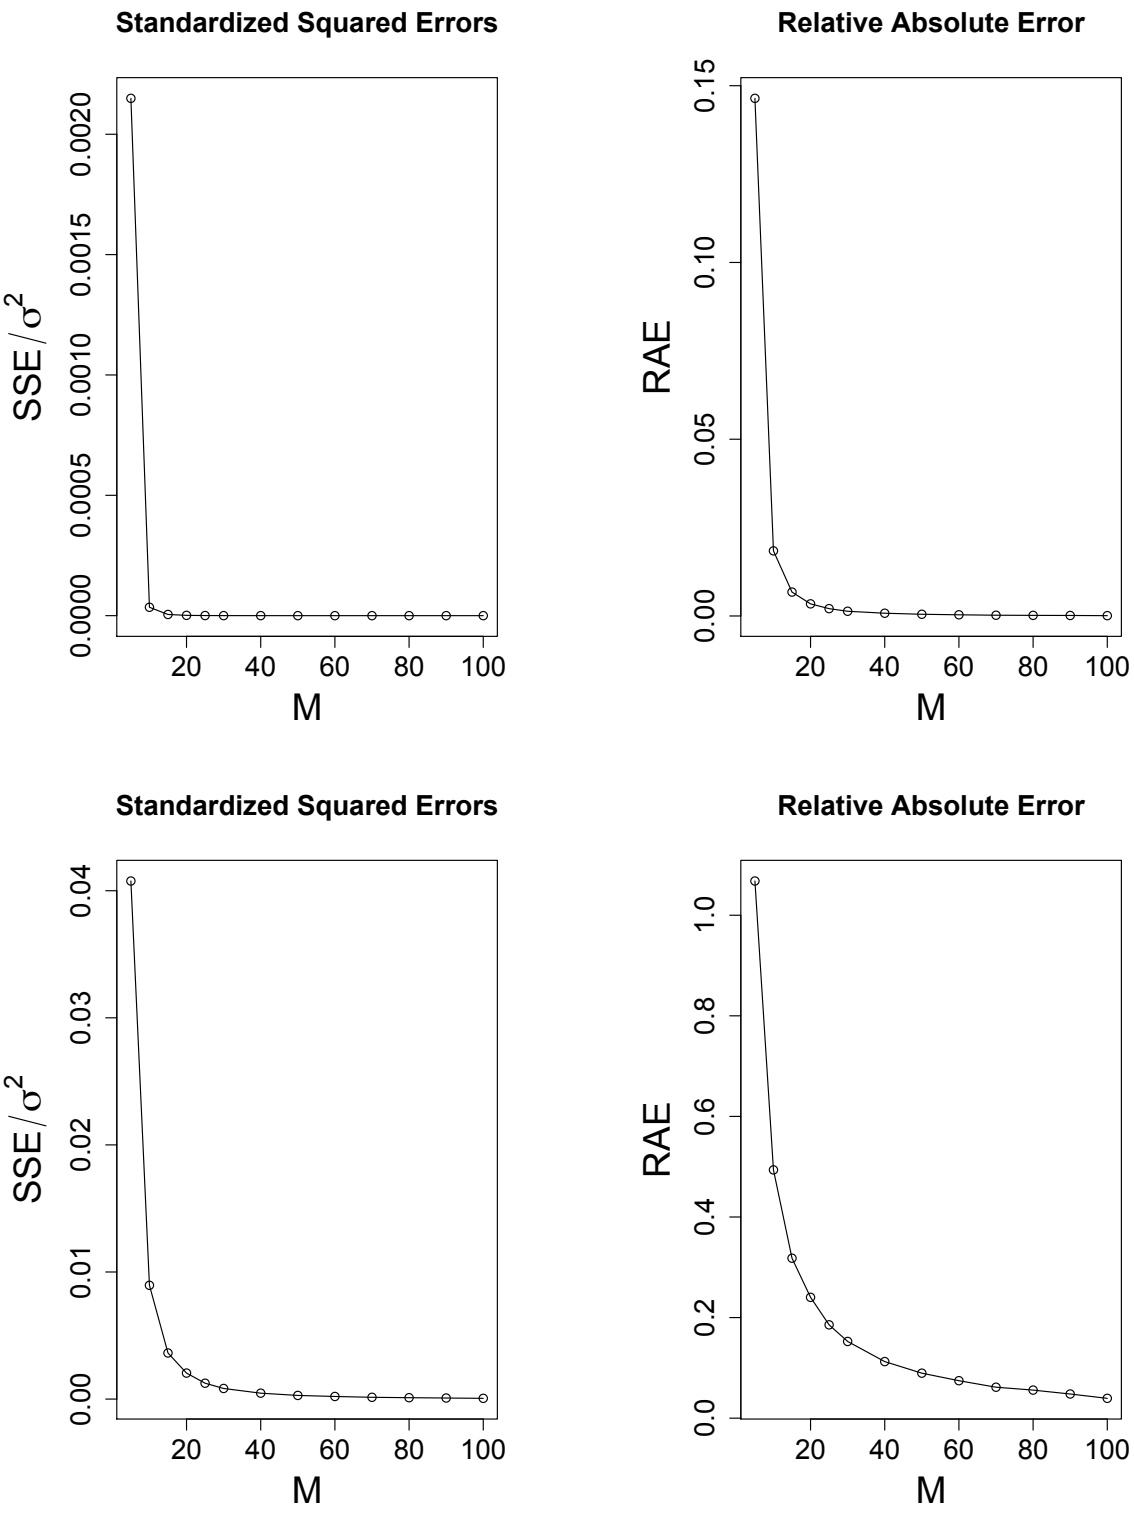

Figure E1: Standardized squared errors and relative absolute errors for smooth (**top**) and non-smooth (**bottom**) integrands. The errors are small in magnitude, particularly in the smooth case, and decay quickly for $M > 20$.